NOKIA

7450 Ethernet Service Switch
7750 Service Router
7950 Extensible RoutingSystem

Advanced Configuration Guide - Part II
Releases Up To 14.0.R5

3HE 11599 AAAA TQZZA 01

Issue: 01

November 2016

Table of Contents

List of tables

List of figures

Preface

About This Guide

The Advanced Configuration Guide is divided into two books, the Part I Guide and the Part II Guide.

Part I provides advanced configurations for basic systems, system management, interface configuration, router configuration, unicast routing protocols, MPLS, services overview, Layer 2 and EVPN services, Layer 3 services, and Quality of Service.

Part II provides advanced configurations for Multi-Service Integrated Service Adapter and Triple Play Service Delivery Architecture.

Parts I and II of the Advanced Configuration Guide supplement the user configuration guides listed below.

The guide is organized alphabetically within each chapter and provides feature and configuration explanations, CLI descriptions and overall solutions. The chapters in the Advanced Configuration Guide are written for and based on several releases, up to 14.0.R5. The Applicability section in each chapter specifies on which release the configuration is based.

Audience

This manual is intended for network administrators who are responsible for configuring the routers. It is assumed that the network administrators have a detailed understanding of networking principles and configurations.

List of Technical Publications

The 7x50 series documentation set also includes the following guides:

- Basic System Configuration Guide

  This guide describes CLI usage, file system management, and boot option file (BOF) configuration, as well as how to configure basic system management, node timing, and synchronization functions.

- System Management Guide

  This guide describes system security features, SNMP and NETCONF features, and event and accounting logs. It covers basic tasks such as configuring management access filters, passwords, and user profiles.

- Interface Configuration Guide

  This guide describes how to provision Input/Output Modules (IOMs), XMA Control Modules (XCMs), Media Dependent Adapters (MDAs), XRS Media Adapters (XMAs), and ports.

- Router Configuration Guide

  This guide describes logical IP routing interfaces and associated attributes such as IP addresses, as well as Redundancy Protocol (VRRP), and Cflowd.

- Unicast Routing Protocols Guide

  This guide provides an overview of unicast routing concepts and provides configuration examples for Routing Information Protocol (RIP), Open Shortest Path First (OSPF), Intermediate-system-to-intermediate-system (IS-IS), and Border Gateway Protocol (BGP) routing protocols, and for route policies.

- Multicast Routing Protocols Guide

  This guide provides an overview of multicast routing concepts and provides configuration examples for Internet Group Management Protocol (IGMP), Multicast Listener Discovery (MLD), Protocol Independent Multicast (PIM), Multicast Source Discovery Protocol (MSDP), Multipoint LDP, multicast extensions to BGP, and Multicast Connection Admission Control (MCAC).

- MPLS Guide

  This guide describes how to configure Multiprotocol Label Switching (MPLS), Resource Reservation Protocol (RSVP), Generalized Multiprotocol Label Switching (GMPLS), and Label Distribution Protocol (LDP).

- Services Overview Guide

  This guide provides a general overview of functionality provided by the routers and describes how to configure service parameters such as Service Access Points (SAPs), Service Distribution Points (SDPs), customer information, and user services.

- Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN

  This guide describes Layer 2 service and Ethernet Virtual Private Network (EVPN) functionality and provides examples to configure and implement Virtual Leased Lines (VLLs), Virtual Private LAN Service (VPLS), Provider Backbone Bridging (PBB), and EVPN.

- Layer 3 Services Guide: Internet Enhanced Services and Virtual Private Routed Network Services

  This guide describes Layer 3 service functionality and provides examples to configure and implement Internet Enhanced Services (IES) and Virtual Private Routed Network (VPRN) services.

- Versatile Service Module Guide (for the 7750 SR)

  This guide describes how to configure service parameters for the Versatile Service Module (VSM).

- OAM and Diagnostics Guide

  This guide describes how to configure features such as service mirroring and Lawful Intercept (LI), and how to use the Operations, Administration and Management (OAM) and diagnostics tools.

- Log Events Guide

  This guide provides general information about log events.

- Triple Play Service Delivery Architecture Guide (for the 7450 ESS and 7750 SR)

  This guide describes the Triple Play Service Delivery Architecture (TPSDA) support and provides examples to configure and implement various protocols and services.

- Quality of Service Guide

  This guide describes how to configure Quality of Service (QoS) policy management.

- RADIUS Attributes Reference Guide (for the 7750 SR)

  This guide describes all supported RADIUS Authentication, Authorization, and Accounting attributes.

- Multiservice Integrated Service Adapter Guide (for the 7450 ESS and 7750 SR)

  This guide describes services provided by integrated service adapters, such as Application Assurance, IPSec, ad insertion (ADI), and Network Address Translation (NAT).

- Gx AVPs Reference Guide (for 7750 SR)

  This guide describes Gx Attribute Value Pairs (AVPs).

- Acronyms Reference Guide

  This guide lists acronyms used in the customer documentation for 7750 SR, 7450 ESS, and 7950 XRS.

Multi-Service Integrated Service Adapter

In This Section

This section provides MS-ISA configuration information for the following topics:

Application Assurance — Application Identification and User-Defined Applications

This chapter describes Application Assurance (AA) Application Identification and User-Defined Applications configurations.

Topics in this chapter include:

Applicability

This example is applicable to all 7750, 7450 and 7750-SRc chassis supporting Application Assurance and was tested on release 12.0.R4.

There are no specific pre-requisites for this example.

Overview

This chapter is intended for Application Assurance (AA) network architects and engineers. It provides best practice information to customize the AA policy and classify any type traffic to meet the service provider reporting, charging or control requirements.

In addition to the signatures built and supported by Nokia, service providers can create their own application signatures based on various criteria. This customization capability can be used to classify traffic hosted on the provider network (web portal, streaming service) or hosted on the Internet and not yet covered by the default AA signature set.

Basics and Terminology

The following main components are used for AA classification:

- **Application Filters —** App-filters are used to define applications based on Layer 3 to Layer 7 criteria. They provide a mapping between one or more protocol signatures or customized traffic patterns into an application of interest.
- **Application** — Such as BitTorrent®, Netflix®. Traffic is classified into applications using app-filters.
- **Application Group** — Such as peer-to-peer, multimedia streaming. For the purpose of reporting and control, applications of similar type/function can be grouped together in Application Groups (App-Group).
- **Charging Group** — Such as zero rating, default. For the purpose of charging or control, applications and app-group can be grouped together in charging groups.

The following table is a high level example to illustrate how app-filters are used to defined applications and show their logical grouping into app-group and charging group.

Figure 1 **App-Filters/Applications/AppGroup**

Maximum Flexibility to Identify Standard and Custom Applications of Interest

Criteria	App-Filter (ordered list of entries, ACL like)	Application	Application Group	Charging Group
- Protocol - Expression: (HTTP, SIP, H323, TLS, RTSP) - L4 Server Port - IP Server Address - Flow Direction - Custom Protocol	Expression - http: yahoo.com	Yahoo	Web	CG#1 - Default
	Expression - http: maps.google.com	Google Maps		CG#2 - Zero Rating
	Expression - http: facebook.com	Facebook	Social Networking	
	Protocol: ftp_control, ftp_data	FTP	File Transfer	CG#1 - Default
	Protocol: bittorrent, dht, utp	BitTorrent	Peer to Peer	
	Protocol: emule	Emule		

Flexible classification/identification rules (apps-filters) to identify:
- Standard applications
- Custom defined applications

Flexible applications/app-group creation and mapping for:
- Reporting
- Control (redirect, enrichment, policing...)

Independent charging group mapping for differentiated billing.

al_0680

- BitTorrent® and Emule® applications are defined using their protocol signature and grouped in the P2P app-group.
- FTP application is defined using both ftp_data and ftp_control protocol signatures, the app is mapped in the file transfer app-group.

- Google Maps® and Yahoo® web sites are defined using http expression and grouped together in the Web app-group.

Configuration

Classification Criteria (App-Filter)

The operator can take full advantage of the flexible AA policy configuration to classify traffic from any application of interest using various criteria ranging from Layer 3 to Layer 7 expressions.

Expression match criteria allows to further refine traffic classification by identifying traffic from HTTP, HTTPS (SSL/TLS), SIP, H323, RTSP, Citrix protocol signatures.

The different app-filter match criteria are listed below:

- L7 Expression
  - HTTP: Host, URI, User Agent, Referer
  - SSL/TLS: Certificate Org Name, Common Name, SNI
  - H323: Product-ID
  - SIP: URI, User Agent, Media Type
  - RTSP: Host, URI, User Agent
  - Citrix:Application Published Name
  - RTMP:Page-host, page-uri, swf-host, swf-uri
- IP Protocol Number
- IP Server Address
- TCP/UDP Server Port
- Custom Protocol
- Protocol Signature

The following operators are supported to define expression based app-filters:

^ :	Expression start with
$:	Expression end with
*:	Wildcard - anything before or after
\I:	Forces case sensitivity
\d:	Any single decimal digit [0-9]

\.: Any single character
*: Asterisk character

Examples of expression match combinations:

^abcd*: match ‘abcd’ at beginning, can end with anything
abcd: match ‘abcd’ anywhere
*abcd$: match ‘abcd’ at the end
^abcd$: exact expression match ‘abcd’
^ab*cd$: string starts with 'ab', ends with 'cd' (anything else in between)
^ab\dcd$: string starts with 'ab', followed by a decimal digit, ends with 'cd'

Note: It is possible to combine different criteria or expressions within the same filter in which case an implicit AND operation between the criteria within the same filter is done by the system.

Application Definition Example

The example below provides a basic configuration example with the application FTP made of two protocol signatures ftp_control and ftp_data; the application is mapped into the application group file transfer:

Create the application group.

```
configure application-assurance group 1:1 policy
    app-group "File Transfer"
    exit
```

Create the application.

```
configure application-assurance group 1:1 policy
    application "FTP"
        app-group "File Transfer"
    exit
```

Create the app-filters.

```
configure application-assurance group 1:1 policy
    app-filter
        entry <1..65535> create
            protocol eq "ftp_data"
```

```
            application "FTP"
            no shutdown
        exit
        entry <1..65535> create
            protocol eq "ftp_control"
            application "FTP"
            no shutdown
    exit
```

Note: Once the application is created the operator is expected to configure the collection of statistics at the subscriber level for this new application (usually only for business VPNs).

User-Defined Applications

General Recommendations

In order to classify traffic properly it is recommended to follow the guidelines and best practices defined in this section before creating a new application:

- Analyze the application traffic
  - Identify what type traffic is used (Wireshark®).
  - Use the application the same way the end user would use it, the same application can create various flows.
- Configure the appropriate App-Filters
  - Following the analysis of the application done above, create the application.
  - Follow the App-Filter best practices chapter to understand in which range to add the filters.
  - More than one App-Filters can be required to identify a single application.

AppDB/Default AA Policy

The default AA policy called AppDB (Application Database) is provided by Nokia and should be used on most deployments. Contact your regional support organization for more details on how to obtain it.

This configuration includes applications and application-groups most Providers can use by default and is designed to allow the addition of any custom entries required by Service Providers to identify additional services/applications.

Before adding new entries to the template and customizing the configuration it is recommended to follow the next guidelines on app-filters and ranges. These guidelines are key to allow an easy upgrade path from the policy configuration provided by Nokia.

App-Filters

App-Filters are an ordered list of entries. It is important to keep the order of this list consistent with the classification objective.

For instance a common configuration mistake is to configure a filter rule for the HTTP protocol signature before HTTP expression filters. If that was the case then app-filters using HTTP expressions would not be used as the system would find an acceptable match with the protocol signature before walking the list of expressions configured. This mistake is described in the example below:

```
entry 100 create
  description "Default HTTP Protocol"
  protocol eq "http"
application "HTTP"
  no shutdown
exit
entry 110 create
  description "Google"
  expression 1 http-host eq "*.google.com$"
  application "Google"
  no shutdown
exit
```

This is an incorrect AppFilter order. App-filter entry #100 will always match before the http expression entry #110.

Note: It is not necessary to specify a protocol when defining an expression filter, the protocol is implicit based on the type of expression match criteria used (for instance, http, sip, h323).

App-Filters Ranges

The App-Filter list is an ordered list, it is key to configure each app-filter in the right order and in the proper range.

The operator can customize the policy and create applications and app-filters by using the following ranges shown in Table 1 (other ranges are used by the default policy):

Table 1 **Customer Reserved App-Filter Ranges**

Range Name	Description	Start	End
Top range	Top range, matches before any other filters	1	1499
High priority	Matches before the other filters.	2000	3999
Expression range A	HTTP Host, Host+URI ; optionally with IP/Port match	19000	22999
Expression range B	Other Expression Match ; optionally with IP/ Port match	33000	34999
Extended protocols	Protocol-signature + Port|IP|Dir. match	40000	41999
Custom protocols	Custom protocol signature match	61000	61499
Trusted/validate ports	1st packet validate, 1st packet trusted match	61500	61999

Ordering Basics:

- Layer 7 expressions based filters are located before their parent protocol signature (for example, expression matches on http are located before the http protocol app-filter; the same applies to TLS, SIP, H323, RTSP, Citrix).
- HTTP Host and URI are located before the HTTP referer for accounting accuracy (for example, YouTube® from within Facebook® is classified as YouTube®)
- App-filters combining protocol signatures with Layer 4 port, IP protocol, IP address or flow direction are always located before the protocol signature only filter range.

HTTP

Protocol

HTTP is a client/server protocol using TCP/IP at the transport layer to deliver resources such as HTML files, images, videos and more.

HTTP 1.1 enables HTTP clients to use a persistent connection to a server allowing them to reuse the same TCP session for multiple HTTP transactions. Text, images, video, scripts and other objects can be downloaded individually in different transactions through the same TCP session.

Figure 2 describes a typical persistent HTTP connection between a web client and a server with multiple HTTP transactions within the same TCP session:

Figure 2 **HTTP Persistent Connection**

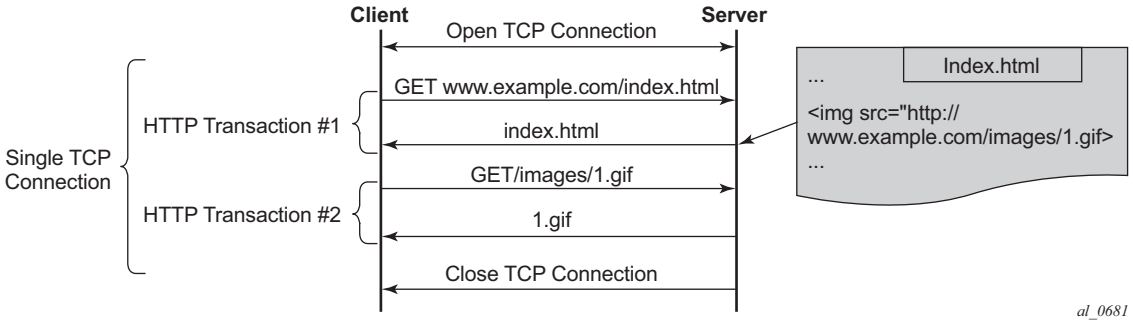

User-defined expression-based HTTP applications will use the first HTTP transaction to classify the flow (optionally this behavior can be modified).

HTTP Request

The example below shows the content of a typical HTTP request to wikipedia.org which includes the following header fields: HTTP Host, HTTP URI, HTTP User Agent and HTTP referer fields:

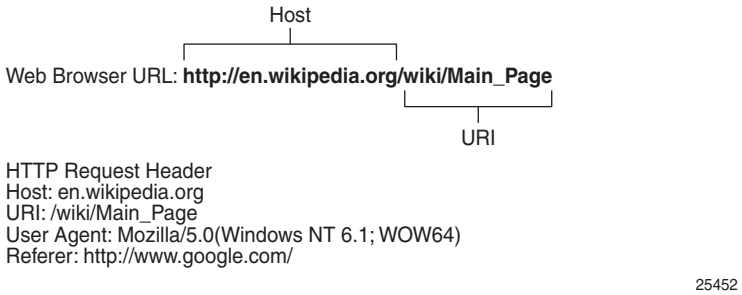

- HTTP Host — Represents the domain name (does not include “http://”).
- HTTP URI — The URL trailer after the host domain name (begins with slash “/”).

- HTTP Referer — The address of the previous web page from which a link to the currently requested page was followed (in this example the referer is www.google.com which means the user clicked on a link from a Google search pointing to wikipedia.org).
- HTTP User Agent — This identifies the web browser or application making the HTTP request.

Configuration Examples

HTTP Host (Wikipedia)

Classifying HTTP traffic from this web site can be done using a single expression tail anchored on the HTTP host:

```
configure application-assurance group 1:1 policy app-filter
    entry <1..65535> create
        description "Wikipedia Web Access" expression 1 http-
host eq "*.wikipedia.org$"
        application "Wikipedia"
        no shutdown
    exit
```

This can be confirmed using Wireshark®.

Figure 3 **Wireshark® www.wikipedia.org**

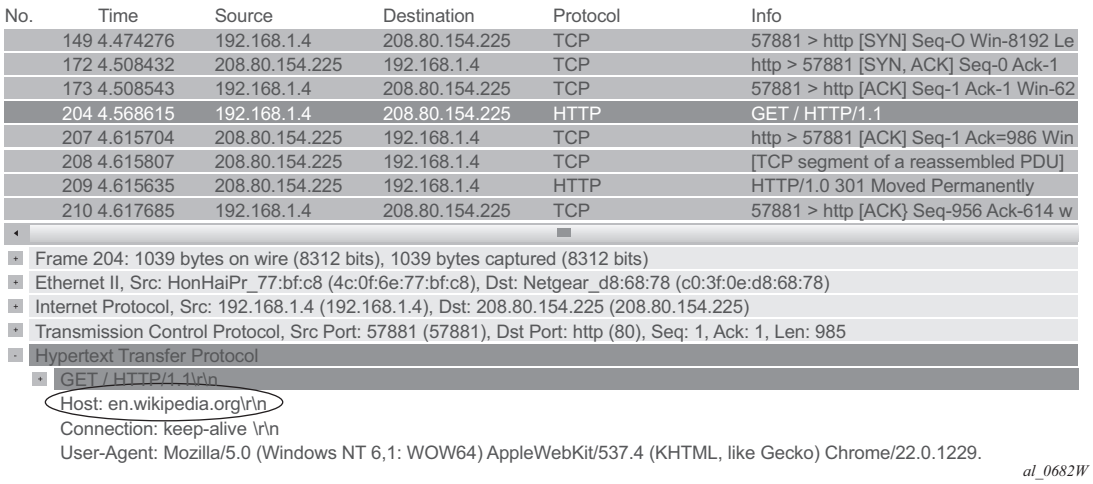

Classification per URI within the Same Host

Operators may need to apply different charging rules to different content located on the same HTTP domain (different URI, same HOST).

Table 2 displays an example of classification rules for the ISP ON-NET content services:

Table 2 **Classification Rules for the ISP ON-NET Content Services**

URL	Charging Rule	AA Application
www.ispdomain.com/video	Rule #1 – 0 Rating	ISP-Portal-Video
www.ispdomain.com/images	Rule #2 – Charge X	ISP-Portal-Images
www.ispdomain.com/*	Rule #3 – Charge Y	ISP-Portal-Default

HTTP 1.1 can reuse the same TCP connection for many transactions to the same server. Classifying each HTTP transaction to www.ispdomain.com independently requires a specific AA configuration.

Prior to SR OS 12.0.R1 the system can be configured to enable traffic classification for all http requests at the AA partition level only therefore affecting all HTTP flows within this partition. SR OS 12.0R1 allows to selectively enable “http-match-all-requests” in app-filters to improve the system performance and limit the HTTP analysis per domain.

The SROS 12.0.R1 configuration example below allows traffic classification of different URIs of the same domain (www.ispdomain.com) independently therefore allowing differentiated charging and control:

- http-match-all-req is enabled on all host+uri app-filters to www.ispdomain.com
- default app-filter required to match any traffic to www.ispdomain.com

```
configure application-assurance group 1:1 policy
    app-filter
        entry <1..65535> create
            description "Zero rated content"
            expression 1 http-host eq "^www.ispdomain.com$"
            expression 2 http-uri eq "^/video*"
            http-match-all-req
            application "ISP Portal Video"
            no shutdown
        exit
        entry <1..65535> create
            description "Image charging"
            expression 1 http-host eq "^www.ispdomain.com$"
            expression 2 http-uri eq "^/images*"
            http-match-all-req
```

```
    application "ISP Portal Images"
    no shutdown
exit
entry <1..65535> create
    description "Default charging"
    expression 1 http-host eq "^www.ispdomain.com$"
    http-match-all-req
    application "ISP Portal Default"
    no shutdown
exit
```

SSL/TLS (HTTPs)

Protocol

HTTPS uses SSL/TLS to encrypt traffic between the client and the server. Since this communication is encrypted it is not possible to identify the HTTP Host or URI. However, AA can still identify the service requested by the subscriber by looking at the TLS certificate information or Server Name Indication exchanged in the clear before the TLS session is established.

Note: SSL/TLS expression based app-filters are not limited to HTTPS. HTTPS is not a protocol in itself, but is HTTP traffic-tunnelled encrypted into SSL/TLS on port 443.

SSL/TLS Certificates

The snapshot (Figure 4) from Wireshark shows the SSL/TLS certificate exchanged using the mobile application **whatsapp**®.

Figure 4 **Wireshark® HTTPS www.whatsapp.com**

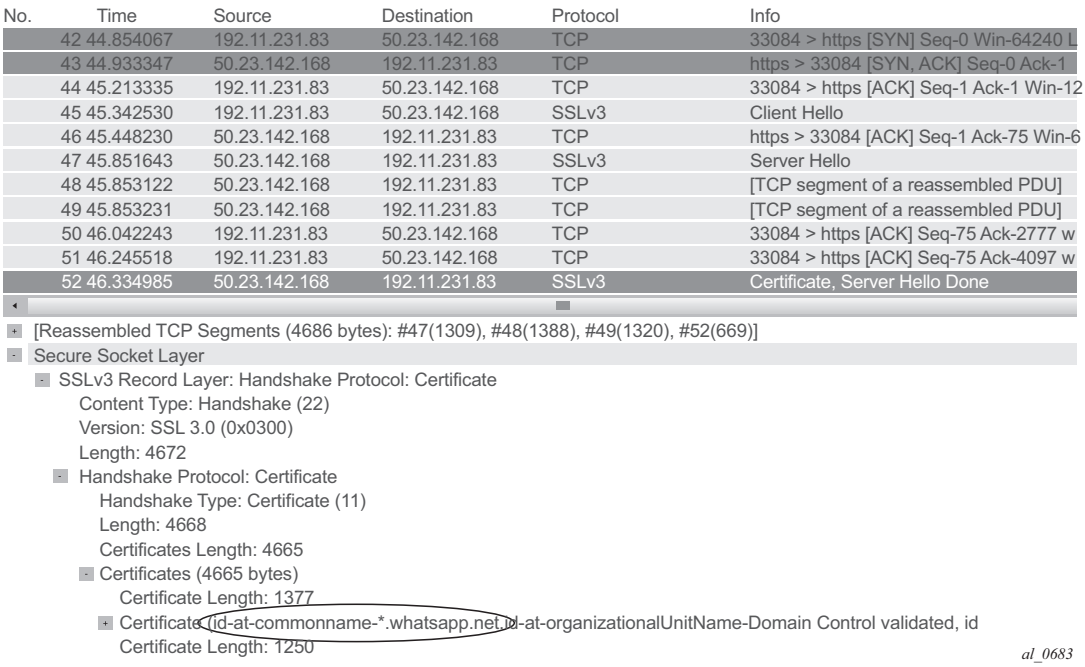

The certificate information can be found in the Server Hello message sent by the server, capturing SSL/TLS (HTTPS) traffic from this application can be done using a single app-filter entry tail anchored on the TLS Common Name Certificate:

```
configure application-assurance group 1:1 policy
    app-filter
        entry <1..65535> create
            description "Whats App tls and image/voice/video traffic"
            expression 1 tls-cert-subj-common-name eq
                         "*.whatsapp.net$"
            application "Whats App"
            no shutdown
          exit
```

Server Name Indication

SSL/TLS traffic can optionally be identified using the Server Name Indication (SNI) which is an extension to the TLS protocol.

The SNI is found in the TLS Client Hello, the http-host expression in the app-filter is reused to classify this traffic:

Figure 5 **HTTPS SNI**

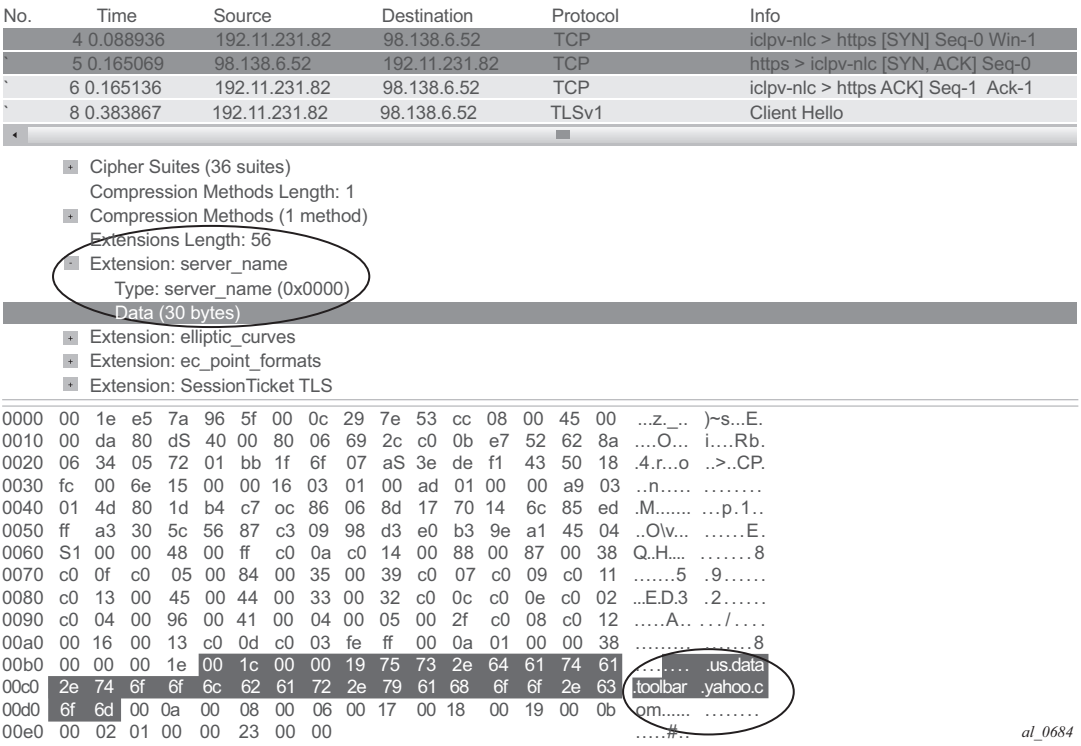


```
configure application-assurance group 1:1 policy
    app-filter
        entry <1..65535> create
            description "Yahoo HTTP or TLS SNI"
            expression 1 http-host eq "*.yahoo.com$"
            application "Yahoo"
            no shutdown
          exit
```

SIP

Protocol

SIP is a signaling protocol used for controlling multimedia communication sessions such as voice and video over RTP. AA automatically monitors SIP control flows and associates RTP/RTCP media flows accordingly in the sip_rtp protocol signature.

The operator can use a SIP expression match criteria in app-filter to further refine traffic classification and identify any additional application on top of the default AA policy. This can be particularly useful in business VPNs to identify voice and telepresence applications.

AA supports SIP expression match criteria on SIP URI, SIP user agent and SIP media type. The snapshot below from Wireshark®shows a SIP control exchange using the voice-video application Vonage® followed by the RTP media audio flow; the expression fields that can be matched using AA app-filters are highlighted:

Figure 6 **SIP Wireshark® Capture**

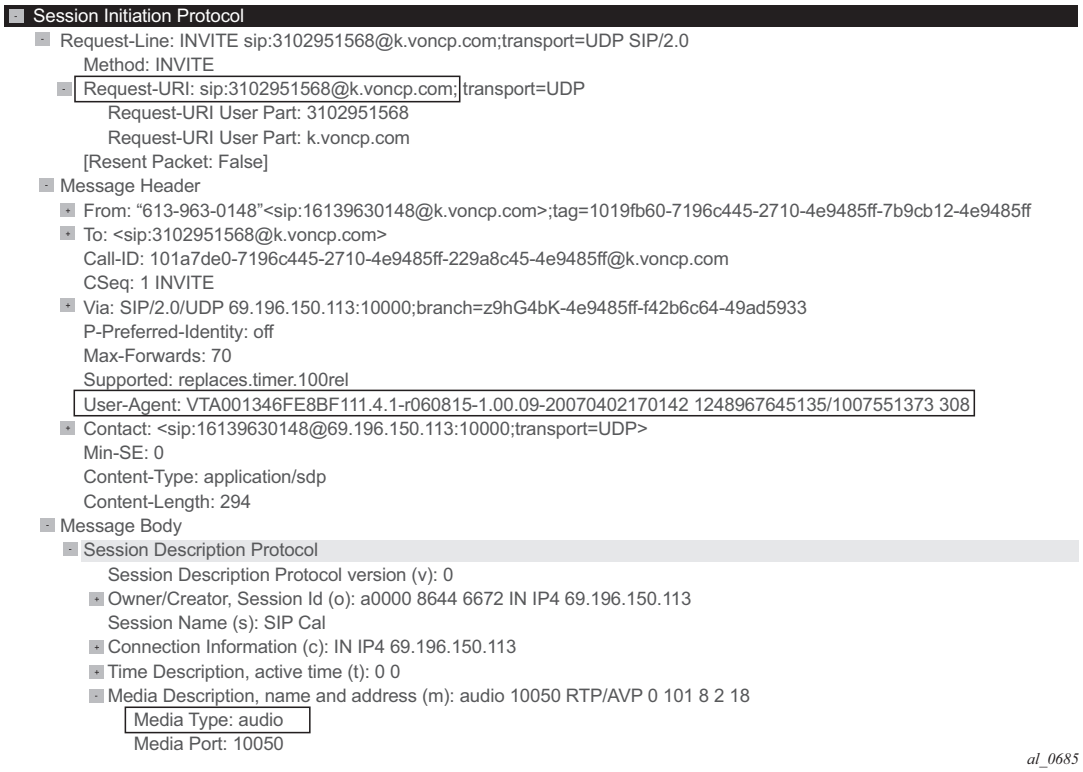

Configuration Example

The configuration example below provides the configuration to classify Vonage® SIP/RTP desktop traffic using SIP URI expression:

```
configure application-assurance group 1:1 policy
    app-filter
        entry <1..65535> create
            description "Vonage"
            expression 1 sip-uri eq "*voncp.com*"
```

```
        application "Vonage"
        no shutdown
    exit
```

H323

Protocol

Similarly to SIP, H323 is a signaling protocol used for controlling multimedia communication sessions such as voice and video over RTP. AA automatically monitors H323 control flows and associates the RTP media flow accordingly in the h323_rtp protocol signature.

The operator can use an H323 expression match criteria app-filter to further refine traffic classification and identify any additional application on top of the default AA policy. This can be particularly useful in business VPNs to identify voice and telepresence applications.

AA supports H323 expression match criteria on the H323 Product ID. The snapshot below from Wireshark shows an H323 control exchange using the Telepresence application LifeSize® followed by the RTP media audio flow; the expression field that can be matched using AA app-filters is highlighted:

Figure 7 **H323 Wireshark® Capture**

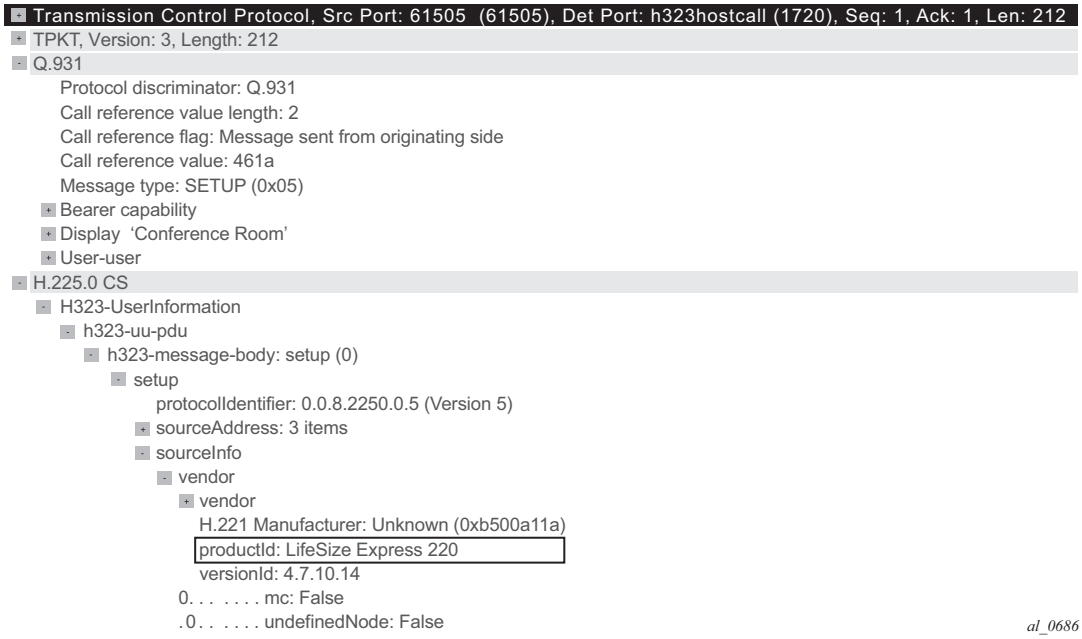

Configuration Example

The configuration example below provides the configuration to classify LifeSize® H323/RTP traffic using the H323 product ID expression:

```
configure application-assurance group 1:1 policy
    app-filter
        entry <1..65535> create
            description "LifeSize H323 traffic"
            expression 1 h323-product-id eq "^LifeSize*"
            application "LifeSize"
            no shutdown
        exit
```

RTSP

Protocol

RTSP is a signaling protocol used for controlling media streaming content such as audio and video over RTP/RDT. AA automatically monitors the RTSP control flows and associates its RTP/RDT media flow with the rtp_rtsp protocol signature.

The operator can use an RTSP expression match criteria app-filter to further refine traffic classification and identify any additional application on top of the default AA policy. This can be particularly useful to identify specific streaming applications.

AA supports RTSP expression match criteria on the RTSP Host, URI, UserAgent. The snapshot below from Wireshark® shows an RTSP setup request to YouTube® followed by the RTP media audio flow; the expression fields that can be matched in RTSP SETUP request using AA app- filters are highlighted:

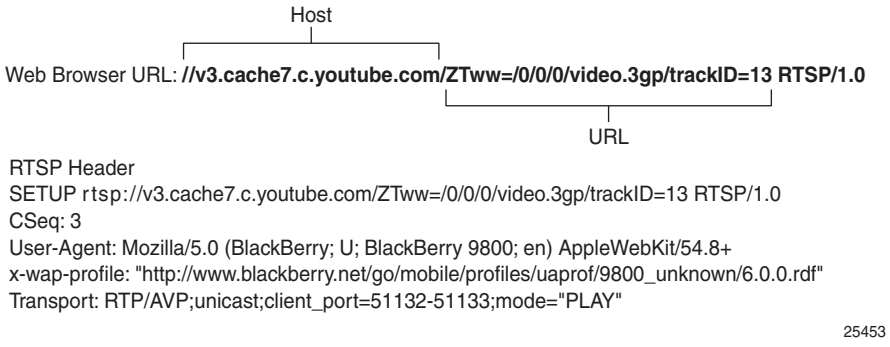

Configuration Example

The configuration example below provides the configuration to classify YouTube® RTSP/RTP traffic using RTSP Host expression:

```
configure application-assurance group 1:1 policy
    app-filter
        entry <1..65535> create
            description "YouTube RTSP/RTP Video"
            expression 1 rtsp-host eq "*.youtube.com$"
            application "YouTube"
            no shutdown
        exit
```

Citrix

Protocol

Independent Computing Architecture (ICA) is a Citrix Systems® protocol used in Citrix's WinFrame, Citrix XenApp (formerly called MetaFrame/Presentation Server), and Citrix XenDesktop products.

Citrix makes it possible to run applications remotely on large servers, thus making better use of server resources while at the same time allowing people using other platforms to use the applications, for example, run Microsoft® Word on a UNIX workstation.

Citrix_ica protocol signature will detect any remote application using Citrix (the protocol needs to be unencrypted and configured to non-seamless). The Citrix ICA session is started from a client and can be anything from Remote Desktop, SAP to Microsoft® Word.

The Citrix expression match app-filter is used to classify traffic based on the Citrix-published application. This published application is configured on the server and in the example above can be for instance RDP, SAP, Word, XLS or Microsoft® Word depending how the server is configured.

Configuration Example

```
configure application-assurance group 1:1 policy
    app-filter
        entry <1..65535> create
            description "Citrix SAP Application"
            expression 1 citrix-app eq "SAP"
            application "Citrix SAP"
            no shutdown
        exit
```

IP Address and TCP/UDP Port

Traffic from specific server(s) can be classified using IPv4/v6 server-address app-filter rules. It is used usually to identify traffic from an internal (on-net) server as opposed to an Internet (off-net) server.

The server-address app-filter automatically detects the client from the server by identifying which side opens the connection. It implicitly classifies traffic based on the server IP address or Port number. For example, if A initiates a TCP connection to B, then flows A->B and B<-A can be classified with a match on server-address = B. Similarly a flow initiated from B to A would be classified using a match on server-address = A.

Server Address

The configuration example below uses a server-address app-filter to classify traffic from server 10.1.1.1 in the application called Application-1:

```
configure application-assurance group 1:1 policy
    app-filter
        entry <1..65535> create
            description "Server #1 10.0.0.1"
            server-address eq 10.0.0.1/32
            application "Application-1"
            no shutdown
        exit
```

Server-Address + Server Port

The configuration example below uses server-address and server-port app-filters to classify traffic from server 10.0.0.2 on port 1234 in the application called Application-2. It is particularly useful when the same server is used to provide different services that need to be classified separately:

```
configure application-assurance group 1:1 policy
        entry <1..65535> create
            description "Server #2 10.0.0.2 port 1234 Only"
            server-address eq 10.0.0.2/32
            server-port eq 1234
            application "Application-2"
            no shutdown
        exit
```

Server Port and Protocol Signature

It is possible to combine a protocol signature with a port number in the same app-filter, this is typically done in business VPNs for specific internal applications not detected using existing AA protocol signatures.

The configuration below classifies a business VPN application running on TCP port 4000 and not detected by any other signatures. It combines the protocol signature unknown_tcp with the desired port number. This allows keeping the classification untouched for the rest of the protocols/applications and is the recommended approach:

```
configure application-assurance group 1:1 policy
    app-filter
        entry <1..65535> create
            description "Business VPN Application X Port 4000"
            server-port eq 4000
            protocol eq unknown_tcp
            application "Busines VPN Application X"
            no shutdown
        exit
```

Note: It is important to follow the app-filter range recommendations for a proper classification of traffic using IP address or port number.

Flow Setup Direction

Traffic can be classified based on flow-setup-direction app-filter. The flow setup direction can be either subscriber-to-network or network-to-subscriber.

Network side and subscriber side is AA terminology related to where AA is enabled:

- In broadband and mobile networks, AA is enabled per subscriber. This means the subscriber side represents the ESM/mobile/transit subscriber while the network side represents Internet or other subscribers.
- In business VPNs, AA is enabled on a VPN SAP/spoke SDP and the subscriber side represents the local VPN site (SAP/spoke/transit).

The example below shows the configuration to classify http traffic hosted by AA subscribers (for example, broadband subscribers running a web server):

```
configure application-assurance group 1:1 policy
    app-filter
        entry <1..65535> create
            description "HTTP Server on the subscriber side"
            flow-setup-direction network-to-subscriber
            protocol eq http
            application "HTTP Server"
            no shutdown
        exit
```

IP Protocol

Traffic can be classified using an IP protocol number for non TCP/UDP traffic.

The example below example provides the configuration to classify ICMP IPv4/v6 traffic:

```
configure application-assurance group 1:1 policy
    app-filter
        entry <1..65535> create
            description "ICMP v4"
            protocol eq "non_tcp_udp"
            ip-protocol-num eq icmp
            application "ICMP"
            no shutdown
        exit
        entry <1..65535> create
            description " ICMP v6"
            protocol eq "non_tcp_udp"
            ip-protocol-num eq ipv6-icmp
            application "ICMP"
            no shutdown
        exit
```

Custom Protocol

Custom protocols can be used to classify TCP/UDP applications using hexadecimal string matching (up to 16 hex octets) at a configurable payload offset in the data payload. The expression string length and offset must not exceed 128 bytes.

To illustrate this feature the Solaris® application GoGlobal is used. It provides remote access to a server (similar to VNC®). The snapshot below (Figure 6) from Wireshark® shows a TCP SYN/ACK session establishment followed by the first data exchange:

Figure 8 **Wireshark® GoGlobal**

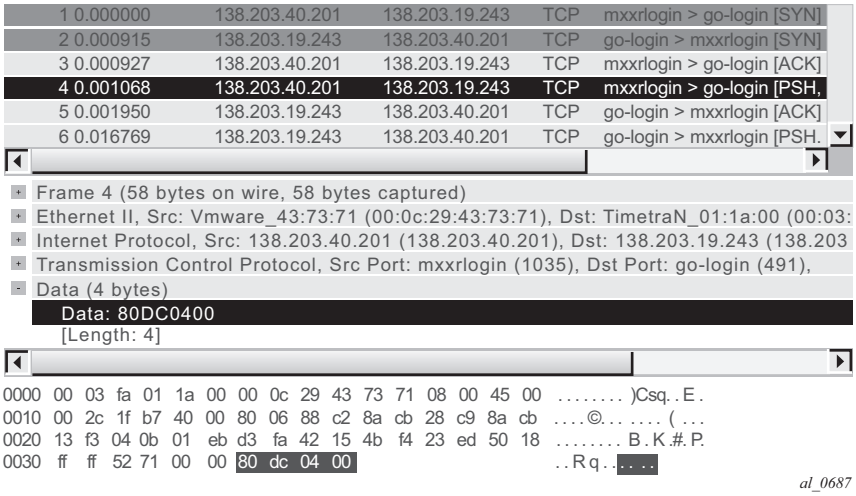

Wireshark® shows that each TCP session payload starts with 80DC0400 (no offset) after the three-way TCP handshake, as a result the configuration required to classify this traffic is described below:

```
configure application-assurance group 1:1 policy
    custom-protocol 1 ip-protocol-num tcp create
        description "goglobal tcp"
        expression 1 eq "\x80\xdc\x04\x00" offset 0 direction client-to-server
        no shutdown
    exit
    app-filter
        entry <1..65535> create
            description "GoGlobal "
            protocol eq "custom_01"
            application "GoGlobal"
            no shutdown
        exit
```

Typical Configuration Mistakes

An operator creating new user-defined applications can make a few typical mistakes which are listed below:

- App-filters in shutdown state — The default app-filter state is shutdown. A n**o shutdown** command must be executed in order for it to be enabled.
- App-filters with no match criteria — This is a more troublesome mistake as it will catch all the traffic entering the filter in a particular application.

Troubleshooting Application Identification

Show Commands

Router/Partition Statistics

Partition level statistics are not updated in real time. Instead, statistics for a particular flow are updated either at flow closure or every five minutes. The fiveminute sliding window interval is a common interval for all flows in a given ISA MDA. Different ISA MDAs will have a different five minute windows as this interval is set at the MDA boot time.

The following command can be used to view the statistics for all applications configured in the ISA Group 1, Partition 1:

```
show application-assurance group 1:1 application count
```

Alternatively it is possible to sort the display by octets, packets, flows:

```
show application-
assurance group 1:1 application count top [octets|packets|flows] [max-count <max-
count>]
```

The operator can also identify which app-filters are being hit by the AA policy per partition (this command is not available per subscriber), it is particularly useful to identify which filters are used and optionally prune unnecessary app-filters from user-defined applications:

```
show application-assurance group 1:1 policy app-filter
```

Note: The app-filter policy is usually relatively large, in which case additional 7x50 SR CLI functionality can be used to filter out the output and only show the relevant information. The example below was created for the application FTP:

```
A:PE# show application-assurance group 1:1 policy app-
filter | match "application \"FTP\""
                                  pre-lines 3 post-lines 2
                    exit
                    entry 44300 create (2 flows, 1205 B)
                        protocol eq "ftp_control"
                        application "FTP"
```

```
        no shutdown
    exit
    entry 44301 create (2 flows, 1401 B)
        protocol eq "ftp_data"
        application "FTP"
        no shutdown
    exit
```

Because partition level statistics are not updated in real time it is recommended for troubleshooting purposes to use subscriber statistics or sub-study statistics.

Subscriber Statistics

Subscriber level statistics can be updated in real time. AA is usually configured by the Operator to collect subscriber level statistics for all application groups in residential and Wifi, while business VPNs typically collect Application group and all applications for each site with AA enabled.

The commands below can be used to view per subscriber statistics for all app-groups or applications configured in ISA Group 1, Partition 1 for the ESM subscriber "Bob" or business VPN SAP 1/1/1:10:

```
show application-assurance group 1:1 aa-sub esm "bob" app-group count
show application-assurance group 1:1 aa-sub sap 1/1/1:10 application count
```

In case only app-group statistics are collected per subscriber, the aa-sub-study feature can be used to collect per application level statistics for selected subscribers, see configuration example below:

```
A:PE# configure application-assurance group 1:1 statistics aa-sub-study application
A:PE>config>app-assure>group>statistics>aa-sub-study# aa-sub esm "bob"
```

Once done, the system will show all application level statistics for this subscriber:

```
show application-assurance group 1:1 aa-sub-study esm "bob" application count
```

Similarly to partition level statistics, aa-sub and aa-sub-study statistics can be sorted by octets, packets, flows:

```
show application-assurance group 1:1 aa-sub-
study esm "bob" application count top [octets|packets|flows] [max-count <max-count>]
```

Note: When the number of flows per ISA card reaches a threshold then per subscriber statistics are not available in real time anymore and only the snapshot command can be used to display the statistics recorded in the previous 5 minute interval window:

```
show application-assurance group 1:1 aa-sub-study esm "bob" snapshot
application count
```

AppFilterMiss

The default policy configuration provides a failsafe application at the very end of the app-filter list to classify any remaining traffic in the AppFilterMiss application. There should never be any traffic in this application. This failsafe filter is used as a debug to make sure that there are no major issues in the configuration.

Note: Traffic can typically be classified as AppFilterMiss when not all protocol signatures are mapped to a particular application. This could happen when upgrading to a new ISA software and enabling new protocol signature detection while not ensuring first that the correct application was provisioned. See the 7x50 SR Release Note upgrade section for more details on AA signature upgrade.

Tools

Flow-Record-Search

Traditional show commands may not provide enough information when troubleshooting flow identification and the operator can use the ISA flow-record-search tool to dump the ISA flow table for more information. This feature comes with a large number of filtering options documented in the user guide.

Each flow gives visibility into: Flow ID, Sub-Type, Sub-Name, Initiator, Direction, Source IP, Dest. IP, IP Protocol, Source Port, Dest. Port, FC, DSCP, Classified, Protocol, Application, App- Group, Charging Group, Packets tx, Bytes Tx, Packets-discarded, Bytes-discarded etc.

See below for the most commonly used commands.

To show all the flows in a given ISA card per ISA group:partition (can be a very long output, up to 3M entries):

```
tools dump application-assurance group 1:1 flow-record-search isa 1/2
```

To show all the flows per AA subscriber in a given group:partition:

```
tools dump application-assurance group 1:1 flow-record-search aa-sub esm "bob"
```

To show all the active flows per AA subscriber in a given group:partition:

```
tools dump application-assurance group 1:1 flow-record-search aa-sub esm "bob" flow-
 status active
```

The flow-record-search command is also available with additional details by adding search-type detail at the end of the command line. Note that due to the length of the output it is recommended to paste the CLI output content in a notepad file.

HTTP Host Recorder

SR OS 11.0.R1 introduced the comprehensive cflowd feature which allows AA to export the HTTP domain extracted from HTTP flows to the 5670 RAM reporting solution. As such, it is the preferred functionality to understand which HTTP Hosts are visible in the network.

Prior to SR OS 11.0.R1, the HTTP host recorder is a tool function available in the 7750 to record HTTP Hosts seen by AA. See the 7750 SR OS Multi-Service Integrated Services Adapter Guide for more details.

```
A:PE# show debug
debug
    application-assurance
        group 1:1
            http-host-recorder
                filter
                    default-filter-action record
                    record http-host-app-filter-candidates
                exit
                rate 100
                no shutdown
            exit
        exit
    exit
exit

A:PE# tools dump application-assurance group 1:1 http-host-recorder top bytes
```

Port Recorder

This function is particularly useful in business VPN (it can also be used in residential networks). The port-recorder AA tool function is similar to the http-recorder. It allows the operator to record which ports are used on selected applications.

It is most commonly used with the applications Unidentified TCP and Unidentified UDP but it can be configured to record any other applications:

```
A:PE# show debug
debug
    application-assurance
        group 1:1
            port-recorder
                application "Unidentified TCP"
                application "Unidentified UDP"
                rate 100
                shutdown
            exit
        exit
    exit
exit

A:PE# tools dump application-assurance group 1:1 port-recorder top bytes
```

Conclusion

This example, which is intended for Application Assurance (AA) network architects and engineers, provides the information required to modify an existing AA policy following AA best practices and guidelines, and provides the necessary troubleshooting information to better understand application classification using Application Assurance.

Application Assurance — App-Profile, ASO and Control Policies

This chapter provides information about Application Assurance (AA) app-profile, Application Service Options (ASOs) and control policy configurations.

Topics in this chapter include:

Applicability

This example is applicable to all 7750, 7450 and 7750-SRc chassis supporting Application Assurance and was tested on SR OS release 12.0.R4.

It is recommended to use the AppDB prior to configuring traffic control policies. The AppDB is a default configuration file to define all of the applications of interest, including all of the relevant application-groups, applications and app-filters to classify traffic, and can be obtained through Nokia's support organization.

Overview

In addition to providing valuable traffic analysis and statistics information using the 7750 Service Router (SR) or 7450 Ethernet Service Switch (ESS) and Application Assurance (AA), one of the key objectives of the AA solution is to provide the tools to manage subscriber traffic at the application level. Examples of traffic management actions include:

- Throttling low priority bandwidth hungry applications during peak hours.
- Prioritizing and remarking selected applications.
- Implementing a walled-garden environment providing open access to selected free web services only, redirecting all other requests from unregistered subscribers to a registration portal with payment services.

- Enrich HTTP Header with subscriber identification parameters to offer subscribers transparent access to premium content.
- In browser notification which triggers the display of administrative, informational or promotional messages in selected browser-sessions.
- Stateful session filtering with Application Level Gateway (ALG) support to protect subscribers against unsolicited flows.
- Parental control services interworking with an external Internet Content Adaptation Protocol (ICAP) server for rating the requested web sites.

Application traffic control policies can be applied as global policies for all subscribers, or they can be activated for individual subscribers or groups of subscribers.

This example describes the basics of activating Application Assurance on a given subscriber through the use of App-Profile and demonstrates the use of static or dynamic traffic control policies using Application Service Options (ASOs) and Application QoS Policies (AQP). It also provides detailed information for configuring Bandwidth, Flow-Count and Flow-Rate Policing including Time of Day (ToD) policing. Other policy control actions can be found in the Advanced Configuration Guide or in the MS-ISA User Guide.

Configuration

Activation of AA Services

App-Profile

Application profiles (app-profile) enable application assurance services for a given Enhanced Subscriber Management (ESM), Distributed Subscriber Management (DSM), or transit subscriber, or for a SAP or spoke SDP which are commonly referred to as **AA-subscribers** (**AA-sub)**. Each app-profile is unique in the system and defines the services that the AA subscriber will receive.

Assigning an app-profile to an ESM subscriber affects every host of that subscriber. Similarly, applying an app-profile to a SAP/spoke SDP will affect all traffic within that SAP/spoke SDP.

App-profiles are defined at the AA group partition level (in case of a partitioned ISA-AA group), see the configuration example below:

```
A:BNG# configure
    application-assurance group 1:1 policy
        app-profile "1-1/15M" create
            description "App-Profile Description"
            divert
            characteristic "Parental Control" value "enabled"
            capacity-cost 15
        exit
```

The app-profile parameters are:

- **divert** — Diverts all traffic from and to this subscriber to an ISA-AA. Configuring **no divert** effectively disables all AA services for subscribers using this app-profile.

  Default value: **no divert**.
- **characteristic** [<*characteristic-name*> **value** <*value-name*>] — one or more optional ASO service characteristics can be used to apply an AA control policy to the subscriber.
- **capacity-cost** <*cost*> — An application profile capacity cost is used to load balance AA subscribers across multiple ISA-AA cards. A common practice is to define a cost proportional to the expected peak BW for the subscribers using this profile (in Kbps or Mbps). The capacity cost is out of the scope of this example. The range is 1 to 65535, default 1.

This app-profile example uses the following naming convention:

<*group-id*>-<*partition-id*>/<*BW*>M where

- <*group-id*> — The ISA-AA group ID on which this profile is created.
- <*partition-id*> — The AA partition ID on which this profile is created.
- <*BW-label*> — Defines the maximum bandwidth used by the subscriber, which is used for aa-subscriber cost load balancing and subscriber rate limiting. The **M** stands for Mbps.

In general the operator can choose to use either ASO characteristics override or multiple app-profiles to apply different AA QoS policies to ESM Subscribers or Business VPN sites. For flexibility and scale it is recommended to use ASO overrides whenever possible. This is described in more details below.

Note: Prior to using special characters in a policy object name the operator should verify the list of special characters supported by the 5620 SAM; for instance the 5620 SAM does not support the use of ":" in the app-profile name therefore it should be avoided.

Residential and Wi-Fi Services

The app-profile can be assigned or modified for ESM, DSM or Transit IP subscribers either at subscriber creation time or while the subscriber is in service:

- Subscriber creation — An app-profile can be assigned at subscriber creation time through RADIUS, DHCP Option 82, Local User Database, static configuration or through a default app-profile.
- In service app-profile modification — An app-profile can be dynamically modified in service through a RADIUS Change of Authorization (CoA). From software release 12.0.R1 an app-profile can also be dynamically modified in service through Gx.

In case no app-profile is returned at subscriber creation by RADIUS, LUDB or DHCP, or when no static configuration is present, the system can apply a default app-profile if configured within the subscriber group-interface (or MSAP policy) sub-sla-mgmt:

```
sub-sla-mgmt
 def-app-profile "1-1/15M"
exit
```

Business VPN and other Service Interfaces

App-profiles are statically assigned to a given SAP, spoke SDP or transit prefix VPN site via the 5620 SAM or CLI.

The following configuration shows how to enable application assurance on a SAP or spoke SDP in a business VPRN service:

```
A:PE>config>service# vprn 100 customer 1 create
 description "L3 Service Customer 1"
 interface "to-site1" create
    address 192.168.1.1/24
    sap 1/1/10:11 create
        app-profile "1-1/15M"
    exit
interface "to-site2" create
    address 192.168.2.1/24
    spoke-sdp 12:100 create
        app-profile "1-1/15M"
    exit
no shutdown
```

Defining Application Service Options

ASOs for Traffic Control - Introduction

To determine which application control policies need to be applied to a AA-subscriber, an app-profile with a number of service characteristics (ASOs) is associated with each subscriber. These service characteristics are then used as match criteria in AQP policy rules to determine which rules to apply.

Therefore ASOs are service characteristics assigned to a subscriber and are used to identify the traffic control policy rule (AQP) applicable to a subscriber or a group of subscribers.

Most policy rules will be applicable to multiple subscriber profiles; nevertheless it is possible that a specific subscriber requires a dedicated policy.

ASO Characteristics and Values

For each service option that can be used by one or more subscribers, an ASO characteristic should be defined with a number of values that represent all available choices for that service characteristic. The names and values of the ASO characteristics are configurable string values; best practice is to use strings that provide a meaningful description of the service characteristic they represent.

Each ASO characteristic requires a default value and each app-profile inherits the default value of all the ASO characteristics created in a given partition unless a characteristic is referenced directly in the app-profile or overwritten as described below.

ASOs are defined at the AA group partition level (in case of a partitioned ISA-AA group). In the configuration example below two different ASO characteristics are defined: “Parental Control” and “P2P-Sub-DL”:

```
BNG>config>app-assure# group 1:1 policy
app-service-options
   characteristic "Parental Control" create
     value "disabled"
     value "enabled"
     default-value "disabled"
   exit
   characteristic "P2P-Sub-DL" create
     value "500k"
     value "1M"
     value "unlimited"
```

```
    default-value "unlimited"
  exit
```

The ASO values and default value of a characteristic can be displayed using a show command:

```
A:BNG# show application-assurance group 1:1 policy app-service-option "P2P-Sub-DL"
===============================================================================
Application-Assurance Application Service Options
===============================================================================
Characteristic "P2P-Sub-DL"
Value                           Default
-------------------------------------------------------------------------------
1M                              No
500k                            No
unlimited                       Yes
===============================================================================
```

When configuring service characteristics for optional service options, it is recommended to configure a default value which will not trigger any AQP policy action (the default value does not match any AQP match criteria) such that the behavior of existing subscribers and app-profiles will not change until the operator specifically configures or signals a non-default characteristic value for the subscriber or the app-profile. In the example above “Parental Control” “disabled” and “P2P-Sub-DL” “unlimited” would have no corresponding AQP by design; therefore if these particular service options were applied to a subscriber they would not match a QoS policy entry.

How to Specify Service Options for AA Subscribers

ASO Assignment in App-Profile

ASOs can be statically assigned in the app-profile; this type of ASO characteristic assignment is typically reserved to the default service options enabled on a large number of subscribers.

Figure 9 shows an example of AA service definition (ASO and app-profile) for a Gold and Bronze service tier definition with the following characteristics:

- Two app-profiles **Gold** and **Bronze**
- **Gold** app-profile — No specific policy actions or ASO characteristics are configured statically in the app-profile.

- **Bronze** app-profile — A specific ASO characteristic and value is assigned to the profile to limit Peer to Peer download traffic to 1Mbps (this example does not show the app-qos-policy nor policer configuration, this will be described later).

Figure 9 **Service Tier Example using ASO, App-Profile and AQP**

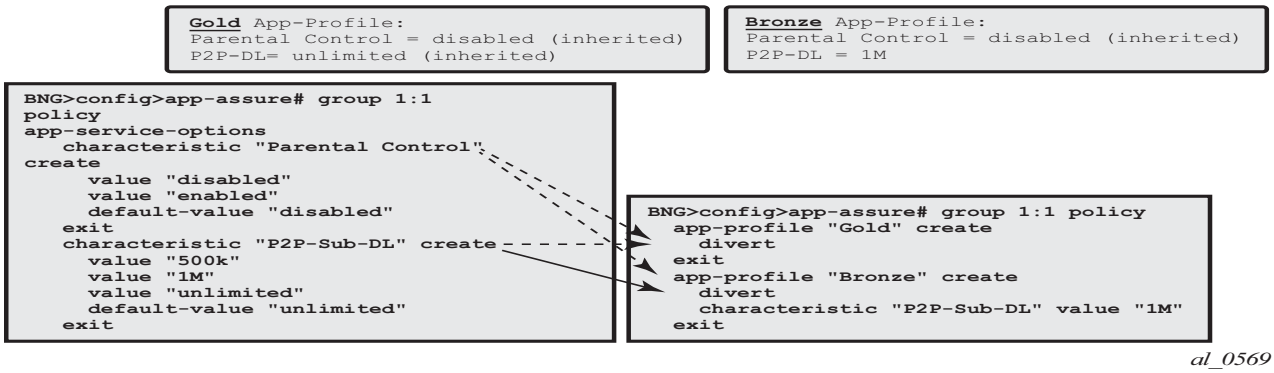

Each app-profile inherits the default values of all the ASO characteristics defined in a AA group-partition; in the example above this is reason why the app-profile Gold inherits “Parental Control” “disabled” and “P2P-Sub-DL” “unlimited”. The app-profile Bronze inherits “Parental Control” “disabled” while “P2P-Sub-DL” “1M” is assigned to this profile statically.

The operator can identify per app-profile which characteristics values are inherited from their default value and which are statically assigned using the following show command:

```
*A:BNG# show application-assurance group 1:1 policy app-profile "Gold"
    app-profile "Gold" create
        divert
        characteristic "P2P-Sub-DL" inherits default-value "unlimited"
        characteristic "Parental Control" inherits default-value "disabled"
    exit

A:BNG# show application-assurance group 1:1 policy app-profile "Bronze"
    app-profile "Bronze" create
        divert
        characteristic "P2P-Sub-DL" value "1M"
        characteristic "Parental Control" inherits default-value "disabled"
    exit
```

Note: Using ASO overrides, described later, it is possible to implement the same choice of AA service options using a single app-profile.

ASO Overrides per Subscriber via RADIUS or Gx

Prior to SR OS 12.0.R1 the operator can assign (and modify: CoA) the app-profile per ESM or Transit-IP subscribers using the “Alc-App-Prof-Str” [26-6527-45] RADIUS attribute.

SR OS 12.0.R1 added support for ASO characteristic overrides for ESM and Transit-IP subscribers via RADIUS using the attribute “Alc-AA-App-Service-Options” [26-6527-193].This attribute can be returned during the subscriber creation process or while the subscriber is in service through RADIUS CoA. Refer to SR OS 12.0 RADIUS Attributes Reference Guide for more details related to the use of the AA RADIUS attributes.

An example of a RADIUS CoA message returned to the system to modify both the app-profile and one ASO characteristic is provided below:

```
NAS-Port-Id = "1/1/5:4088"
Framed-IP-Address = 192.168.211.30
Alc-App-Prof-Str = "1-1/15M"
Alc-AA-App-Service-Options = "P2P-Sub-DL=1M"
```

The ASO characteristics and values assigned to a given subscriber (statically via app-profile or overridden) can be displayed using the following show command:

```
A:BNG# show application-assurance group 1:1 aa-sub esm "sub1" summary
===============================================================================
Application-Assurance Subscriber Summary (realtime)
===============================================================================
AA-Subscriber           : sub1 (esm)
ISA assigned            : 1/2
App-Profile             : 1-1/15M
App-Profile divert      : Yes
Capacity cost           : 1
Aarp Instance Id        : N/A
HTTP URL Parameters     : (Not Specified)
Last HTTP Notified Time : 2014/08/07 12:07:47
-------------------------------------------------------------------------------
Traffic                        Octets              Packets                Flows
-------------------------------------------------------------------------------
...
...
-------------------------------------------------------------------------------
Application Service Options (ASO)
-------------------------------------------------------------------------------
Characteristic                 Value                            Derived from
-------------------------------------------------------------------------------
P2P-Sub-DL                     1M                               dyn-override
Parental Control               disabled                         default
===============================================================================
```

In the show command output above, the **derived from** field describes how the characteristics and values are assigned to the subscriber:

- app-profile — The characteristic's value statically configured in the app-profile.
- dyn-override — The characteristic's value received from RADIUS or Gx.
- default — The characteristic's default value inherited (not statically configured in the app-profile nor dynamically modified).

SR OS 12.0.R1 also introduced support for signaling the app-profile or ASO characteristics override via Gx, see Application Assurance — App-Profile, ASO and Control Policies for more details.

ASO Overrides for Business VPN and Other Services

Since SR OS 9.0.R1, ASO characteristic override values can be statically assigned to business VPN SAP, spoke SDP and transit prefix subscribers.

The operator can provision the AA policy override parameters, multiple characteristics overrides per AA-sub can be defined per override policy, see the configuration example below:

```
A:BNG>config>app-assure# group 1:1 policy-override
    policy aa-sub sap 1/1/5:210 create
        characteristic "P2P-Sub-DL" value "1M"
        characteristic "Parental Control" value "enabled"
    exit
```

Application Control Policies

App-QoS-Policy (AQP)

App-Profile / ASO / AQP Workflow Summary

App-profiles enable application assurance services for a given AA-subscriber. Each app-profile is unique in the system and defines the service that the AA subscriber will receive.

To determine which control policies need to be applied to an AA-subscriber, a number of service characteristics (ASO) are associated with each AA-subscriber.

As described earlier, these service characteristics can either be configured directly within the app-profile or assigned using overrides and they are then used as match criteria in AQP policy rules to determine which application policy rules to apply.

The app-qos-policy (AQP) is an ordered list of entries defining policy actions for flows diverted to Application Assurance. Each AQP entry is composed of match criteria and action(s).

Flows are evaluated against all entries of the AA QoS policy defined in the AA group partition that the subscriber app-profile belongs to (in case of a partitioned AA group).

Figure 10 provides a configuration example summary with app-profile, ASO, AQP and policers:

Figure 10 **App-Profile, ASO, AQP Workflow Summary**

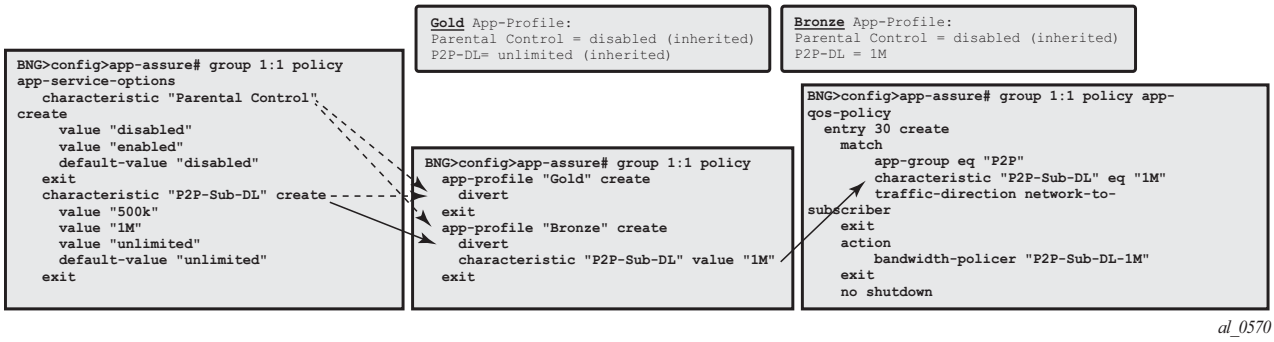

Match and Action Criteria

AQP Match Criteria

Multiple match criteria can be specified per AQP entry in which case the action will only apply to flows that match all criteria. The most common match criteria are: characteristic, application, app-group and charging-group.

The following AA match criteria can be used in an AQP:

- **app-group** {**eq** | **neq**} <*app-group name*>
- **application** {**eq** | **neq**} <*app name*>
- **charging-group** {**eq** | **neq**} <*charging-group-name*>

- **traffic-direction** {**subscriber-to-network**|**network-to-subscriber**|**both**}
- **characteristic** <*characteristic-name*> <**eq**> <*value-name*>: up to 4 characteristics and values per AQP
- **ip-protocol-num** {**eq** | **neq**} <*protocol-id*>
- **src-ip** {**eq** | **neq**} <*ip-address*> or **ip-prefix-list** <*ip-prefix-list-name*>
- **dst-ip** {**eq** | **neq**} <*ip-address*> or **ip-prefix-list** <*ip-prefix-list-name*>
- **src-port** {**eq** | **neq**} <*port-num*> or **range** <*start-port-num*><*end-port-num*>
- **dst-port** {**eq** | **neq**} <*port-num*> or range <*start-port-num*><*end-port-num*>
- **dscp** {**eq** | **neq**} <*dscp-name*>
- **aa-sub** <*aa-sub-name*>

AQP Actions

The following AA traffic control policies can be specified in an AQP:

- **drop**
- **bandwidth-policer** <*policer-name*>
- **flow-count-limit** <*policer-name*>
- **flow-rate-limit** <*policer-name*>
- **remark dscp in-profile** <*dscp*-name> **out-profile** <*dscp-name*>
- **remark fc** <*fc-name*>
- **remark priority** <*priority-level*>
- **http-error-redirect** <*redirect-name*>
- **http-redirect** <*redirect-name*> **flow-type** <*flow-type*> — Redirect traffic to a landing page
- **mirror-source** [**all-inclusive**] <*mirror-service-id*>
- **session-filter** <*session-filter-name*> — Session filter firewall
- **url-filter** <*url-filter-name*>: category based URL Filtering using ICAP
- **http-notification** <*http-notification-name*>
- Additional drop actions:
  - **error-drop**: configure a drop action for packets cut-through due to IP packet errors (bad IP checksums, tcp/udp port 0, etc.)
  - **overload-drop**: configure a drop action for packets cut-through due to overload
  - **fragment-drop**: configure a drop action for IP fragmented packets

Default Versus Application-Specific AQP Policies

Application QoS Policy

It usually requires the examination of a few packets to identify the protocol/application of a flow. When AQP entries are defined to match on IP header criteria (IP address, IP prefix list, TCP/UDP Port Number, IP Protocol, DSCP) or application criteria (application, App-Group or charging group), the AQP action will only be applied to matching application flows after a flow has been classified as a given application.

Default QoS Policy

If the AQP entry does not include match criteria against application (application, app-group and charging-group) or IP header information (IP address, IP prefix list, TCP/UDP port number, IP protocol, DSCP) then the AQP policy will be applied to all matching flows starting with the first packet of a flow before protocol and application identification is complete. Such AQPs are called default subscriber policies.

For an AQP to be qualified as a default subscriber policy, the match criteria must be limited to any combination of ASO characteristic values, traffic direction and optional AA subscriber name.

AQP match and actions for the default QoS policy and application QoS policy are summarized in Table 3:

Table 3 **Default QoS Policy, Application QoS Policy Table**

Policy	AQP Match	AQP Action
Default QoS	ASO characteristic/values traffic direction aa-sub	Remark FC, DSCP, Priority Bandwidth, flow-count, flow-rate policing Session-filter Url-filter Mirror Error-drop, overload-drop, fragment-drop Drop

Table 3 **Default QoS Policy, Application QoS Policy Table (Continued)**

Policy	AQP Match	AQP Action
Application QoS	ASO characteristic/values traffic direction aa-sub application app-group charging-group IP address, IP Prefix List TCP/UDP Port Number DSCP IP Protocol Number	Remark FC, DSCP, Priority Bandwidth, flow-count, flow-rate policing HTTP Notification HTTP Redirect HTTP Enrichment Mirror Drop

To ensure fair access to the ISA-AA bandwidth and flow resources, it is recommended to configure default AQP policy entries limiting bandwidth and flow resources per AA sub.

Figure 11 shows a default subscriber policy limiting the downstream bandwidth (network-to-subscriber direction) to 25Mbps per subscriber:

Figure 11 **Default Downstream Bandwidth Policing**

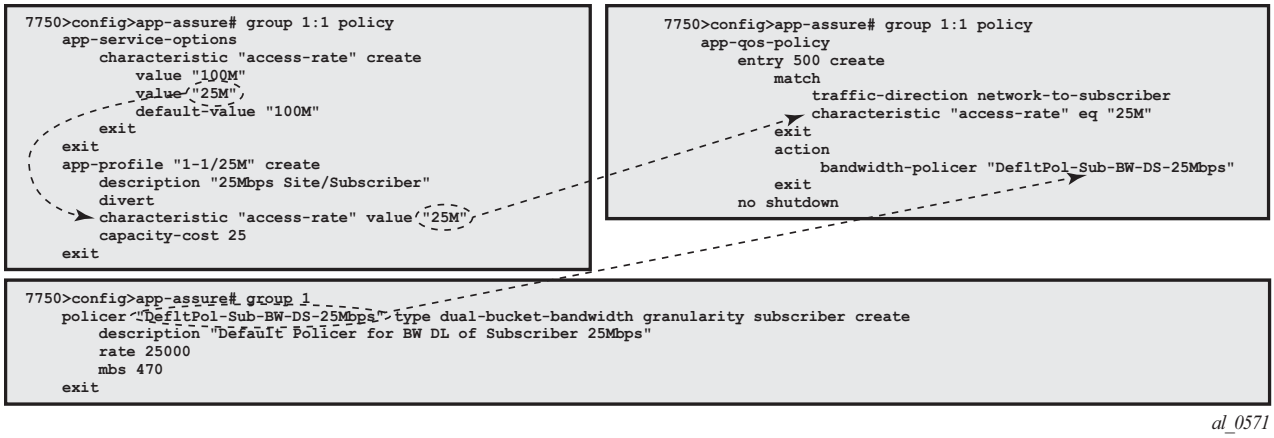

al_0571

Implicit Default Subscriber Policy

Session-filter, url-filter, overload-drop, fragment-drop and error-drop can only be used as part of a default subscriber policy; therefore these actions are not compatible with application or IP header match criteria within the same AQP.

AQP Entries Evaluation

Multiple AQP Match Entries Per Flow

A single flow can match multiple AQP entries, in which case multiple actions can be selected based on the AQP entry's order (the lowest number entry has the highest priority); the drop action takes precedence over any other AQP entry. The maximum numbers of actions that can be applied on a single flow are:

- 1 drop action
- Any combination of (applied only if no drop action is selected)
  - Up to 1 mirror action
  - Up to 1 FC, 1 priority and 1 DSCP remark action
  - Up to 4 BW policers (1 single rate AA-Sub, 1 dual rate AA-Sub, 2 single rate system level)
  - Up to 12 flow policers (3 subscriber flow-count, 3 subscriber flow-rate, 3 system flow-count, 3 system flow-rate)
  - Up to 1 HTTP Redirect
  - Up to 1 HTTP Error Redirect
  - Up to 1 HTTP Enrichment
  - Up to 1 URL-Filter
  - Up to 1 HTTP-Notification
  - Up to 1 Session-Filter Firewall
- 1 error-drop
- 1 overload-drop
- 1 fragment-drop

An AQP entry match that would cause the above limits to be exceeded is ignored (no actions from that rule are selected) and the conflict counter for this AQP is incremented.

The operator can display hits and potential conflicts per AQP entry using the following show command:

```
A:BNG# show application-assurance group 1:1 policy app-qos-policy
===============================================================================
Application QOS Policy Table
===============================================================================
Entry         Admin State                    Flow Hits          Flow Conflicts
-------------------------------------------------------------------------------
30            in-service                             0                      0
-------------------------------------------------------------------------------
```

```
No. of AQP entries: 1
==============================================================================
```

AQP Evaluation

Flows are evaluated against all entries of the AA QoS Policy at different steps during the lifetime of the flow:

- **Flow creation** — The default subscriber policy AQP entries for matching flows are applied starting with the first packet of a flow so before application identification completes.
- **Application identification completion**— The application QoS policies are applied once flow identification has been completed.

Note: The default QoS policy entries are applied to the subscriber's flows for packets received before and after application identification is completed.

- **Policy change** — When a configuration change is applied to the AA policy by executing the commit command on the AA group:partition policy, all diverted flows for subscribers using this policy partition will be evaluated again against all AQP entries. This re-evaluation happens as a paced background task; hence AQP control changes may not be applied immediately to all existing flows.

Policing

Policers

AA policer templates are configured as part of the AA Group configuration by specifying the policer name, type and granularity. Policers are unidirectional by definition so that separate policers must be defined per flow direction if the traffic needs to be policed in both directions (a separate AQP for each flow direction is therefore required as well).

The operator can configure the following types of policers:

- Bandwidth Policers
  - Single bucket system level
  - Single bucket AA subscriber level

  - Dual bucket AA subscriber level
- Flow Count Policer: system or AA subscriber level
- Flow Setup-Rate Policer: system or AA subscriber level

Subscriber level policers are instantiated per AA sub, meaning:

- The system automatically uses a dedicated policer for every single subscriber, even when multiple subscribers match the same AQP entry.
- The same policer can be referenced in different AQP entries; in this case all subscribers' flows matching any of these AQP entries are policed by the same subscriber policer. Example: if the same subscriber level policer '1Mbps' is referenced in AQP entry 100 matching application BitTorrent and in AQP entry 110 matching application EDonkey, then the sum of both the BitTorrent and EDonkey traffic cannot exceed 1Mbps.

System level policers on the other hand are shared by all AA subscribers matching a given AQP entry. These policers are typically used in residential and Wi-Fi service deployments to limit the total bandwidth for an application or application group, for all subscribers or for a group of subscribers on the system or partition. An example would be a system level 500Mbps policer to limit the aggregated downstream bandwidth of "Peer to Peer" applications for all subscribers with a "Bronze" app-profile to 500Mbps.

Note: In case multiple ISA-AA cards are used per system, the overall maximum throughput using a system level policer is equal to the policer rate limit times the number of ISA cards in the system.

Bandwidth Policing

Single Bucket Subscriber/System Bandwidth Policer

Single bucket policers police the matching traffic against a configured peak-information-rate (PIR). Traffic above the PIR can be marked as out of profile or dropped.

The configuration template for a single rate bandwidth policer is as follows:

```
BNG>config>app-assure# group 1
    policer <policer-name> type single-bucket-bandwidth
                                granularity {subscriber|system} create
        description <string>
        rate <pir-rate-in-Kbps>
        mbs <max-burst-size-in-Kbytes>
```

```
        adaptation-rule pir {max|min|closest}
        tod-override <tod-override-id>
        action permit-deny|priority-mark
```

where:

- **action** — Defines the action that must be taken by the policer for non-conforming traffic.
- **permit-deny** — Non-conforming packets will be dropped.
- **priority-mark** — Non-conforming traffic will be marked as out of profile (increasing the chances that non-conforming packets will be discarded in case of congestion on the egress queues).
- **rate —** Peak information rate in Kbps.
- **mbs —** Maximum burst size in Kbytes.
- **adaptation-rule pir <max|min|closest>** — The policers work at discrete operational rates supported by the hardware. The adaptation rule specifies how the actual operational policer rate (supported by the hardware) must be selected as compared to the configured PIR. During operation, both the operational and configured rate can be displayed using the operational **show application-assurance group** <*n*> **policer** <*policer-name*> **detail** command.
- **tod-override** — Defines a time of day override policy applicable to a policer, this is described in more detail at the end of the policing section.

A single bucket subscriber level policer configuration example is shown below:

```
BNG>config>app-assure# group 1
    policer "P2P-Sub-DL-1M" type single-bucket-
bandwidth granularity subscriber create
        rate 1000
        mbs 19
    exit
```

A single bucket system level policer configuration example is shown below:

```
BNG>config>app-assure# group 1
    policer "P2P-Sys-DL-100M" type single-bucket-bandwidth granularity system create
        rate 100000
        mbs 1875
    exit
```

Dual Bucket Subscriber Bandwidth Policer

Dual-bucket policers police the matching traffic against a configured peak information rate (PIR) and committed information rate (CIR). Traffic below CIR is marked in profile, traffic between CIR and PIR is marked as out of profile, and traffic above the PIR is dropped.

Dual-bucket policers can only be used as subscriber policers; system policers cannot be defined as dual-bucket policers.

The configuration is similar to the single-bucket policer, but adds the configuration of a CIR and a Committed Burst Size (CBS), and the action cannot be configured:

```
BNG>config>app-assure# group 1
    policer <policer-name> type dual-bucket-bandwidth
                                  granularity {subscriber|system} create
        description <string>
        rate <pir-rate-in-Kbps> cir <cir-rate-in-Kbps>
        mbs <max-burst-size-in-Kbytes>
        cbs <committed-burst-size-in-Kbytes>
        adaptation-rule pir {max|min|closest} cir {max|min|closest}
```

A dual-bucket subscriber level policer configuration example is shown below:

```
BNG>config>app-assure# group 1
    policer "P2P-Sub-DL-2M-DB" type dual-bucket-
bandwidth granularity subscriber create
        rate 2000 cir 1000
        cbs 19
        mbs 38
    exit
```

MBS/CBS Calculation for Bandwidth Policers

The default MBS/CBS value of a bandwidth policer is set to 0. This value can and should be modified by the operator to allow proper interworking with TCP based applications.

The formula to calculate the MBS or CBS buffer size, as documented in RFC 6349, *Framework for TCP Throughput Testing*, is:

Buffer (B) = Rate (bps) / 8 * RTT (s)

For Internet applications it is recommended to use a common Round Trip Time (RTT) of 150 msec.

An example using a single bucket subscriber level policer rate of 10000 Kbps:

MBS (B) = 1,000,000 / 8 * 0.150 = 18750 Bytes or 190 KB.

Note that these policer values may need to be further adjustment depending on the application.

Flow Rate Limit Policer

Flow rate limit policers police the maximum number of new flows that are accepted per second for matching traffic. The configuration is similar to the single-bucket bandwidth policer, with the rate and MBS now expressed in flows/sec and flows, respectively.

```
BNG>config>app-assure# group 1
    policer <policer-name> type flow-rate-
limit granularity {subscriber|system} create
        description <string>
        rate <flow-rate-in-flows/sec>
        mbs <max-burst-size-in-flows>
        adaptation-rule pir  {max|min|closest}
        action permit-deny|priority-mark
```

This type of policer is primarily used for the default subscriber AQP policy in order to limit the maximum number of flow/seconds allocated per AA subscriber.

Note that in case the policer is used as part of the default AA subscriber policy then the **priority-mark** action has the effect to cut-through non conformant traffic in the ISA instead of drop using **permit-deny**.

Flow Count Limit Policer

Flow count limit policers police the maximum number of concurrent flows for matching traffic:

```
BNG>config>app-assure# group 1
    policer <policer-name> type flow-count-
limit granularity {subscriber|system} create
        description <string>
        action permit-deny|priority-mark
        flow-count <max-number-of-flows>
```

This type of policer is primarily used for the default subscriber AQP policy in order to limit the maximum number of concurrent flows allocated per AA subscriber.

Note that the “priority-mark” has the effect to cut-through non conformant traffic in the ISA instead of drop using “permit-deny”.

Time of Day Policing

Software release 11.0.R1 introduced support for time-of-day (ToD) policer override. Up to 8 override rates with time of day specifications can be defined per policer, this time of day override using the system local time.

ToD overrides are supported for all policer types described in the previous section (bandwidth, flow-count, flow-rate) and can be configured using either daily or weekly patterns.

The configuration of ToD override on daily or weekly basis is shown in the following template:

```
BNG>config>app-assure# group 1
    policer "P2P-Sub-DL-1M-TOD" type single-bucket-bandwidth
                                          granularity subscriber create
        action permit-deny
        rate 1000
        mbs 19
        adaptation-rule pir closest
        tod-override <override-id>
            description <string>
            time-range daily start <start-time> end <end-time>
                                    [on <day> [<day>...(upto 7 max)]]
            time-range weekly start <day,start-time> end <day,end-time>
            rate 2000
            mbs 38
```

where:

- **tod-override** <*override-id*> — Up to 8 override-ids (with value 1-255) can be configured per policer.
- **time-range** — Can be configured to be triggered.
  - On a daily basis at the indicated start/end-time on the specified days.
  - On a weekly basis at the indicated start day+time and end-day+time.
  - Times can be indicated as <hh>:<mm> with a 15-minute granularity for the minutes (mm = 0 | 15 | 30 | 45).

A configuration example for a single bucket system level bandwidth policer with the following ToD-override patterns follows:

- Default Rate Limit: 300Mbps
- Rate Limit override to 100Mbps between 5PM and 10PM

- Rate Limit override to 200Mbps between 10PM and 12PM

```
BNG>config>app-assure# group 1
     policer "P2P-Sys-DL-300M-TOD" type single-bucket-bandwidth
                                          granularity system create
        description "Peer to Peer Policer System level Policer"
        rate 300000
        mbs 5625
        tod-override 1 create
            description "Override busy hour #1"
            time-range daily start 17:00 end 22:00
            rate 100000
            mbs 1875
            no shutdown
        exit
        tod-override 2 create
            description "Override busy hour #1"
            time-range daily start 22:00 end 24:00
            rate 200000
            mbs 3750
            no shutdown
        exit
```

The operator can display which policing rate is applied at any moment in time together with all configured override rates using the following command:

```
show application-assurance group <n> policer <policer-name> detail
```

Design and Configuration Examples

Default AA QoS Policy

To ensure fair access for all subscribers to the ISA-AA resources, and avoid that a disproportionate amount of ISA-AA resources are used by one or more subscribers which are misbehaving or receiving large traffic bursts from the Internet, it is recommended to configure the following three types of subscriber-level default AA QoS policies:

- **A default bandwidth policer** to limit the downstream bandwidth per subscriber (upstream bandwidth is already limited by ESM/SAP access ingress IOM QoS).
- **A default flow count policer** to limit the maximum number of active flows per traffic direction per subscriber. The operator can choose to drop or cut-through non conforming traffic.
- A **default flow rate policer** to limit the maximum flow setup rate per traffic direction per subscriber. The operator can choose to drop or cut-through non conforming traffic.

The minimum set of app-profiles used in a network is typically determined by the different access bandwidth rates; services characteristics are then used for each profile to apply a default QoS policy to limit bandwidth and flow resources accordingly.

In theory, it is possible to configure a set of default policers for every individual access bandwidth rate that is offered to a subscriber. This would however result in a large number of policers and corresponding ASO values plus app-profiles that need to be configured. Therefore, a best practice guideline is to define a small number of bandwidth ranges (not more than five to ten) that cover the full offered access bandwidth spectrum, and define for each bandwidth range a default bandwidth policer plus flow policers with appropriate limits.

As an example, assuming a residential deployment with 2 bandwidth ranges of up to 25Mbps and 100Mbps, the configuration below provides:

- Complete ASO and app-profile configuration.
- Default QoS policy for subscribers in the 25Mbps range including bandwidth.
- Flow count and flow rate policers are configured by default as permit-deny. Non conforming traffic is dropped which is common for residential deployments; alternatively the operator can decide to configure these policers as priority-mark to cut-through traffic in the ISA-AA.

In this example the resources are limited per subscriber based on their access rate maximum speed from which flow count and flow rate are derived.

App-Profile and ASO

The configuration below provides the app-profile and ASO characteristics used for the default subscriber AQP policy for the 25Mbps and 100Mbps access bandwidth range:

```
BNG>config>app-assure# group 1:1 policy
    app-service-options
        characteristic "access-rate" create
            value "100M"
            value "25M"
            default-value "100M"
        exit
    exit
    app-profile "1-1/25M" create
        description "25Mbps Site/Subscriber"
        divert
        characteristic "access-rate" value "25M"
        capacity-cost 25
    exit
    app-profile "1-1/100M" create
        description "100Mbps Site/Subscriber"
```

```
        divert
        characteristic "access-rate" value "100M"
        capacity-cost 100
    exit
```

Default Bandwidth Policing – 25Mbps AA-Sub

```
BNG>config>app-assure# group 1
    policer "DefltPol-Sub-BW-DS-25Mbps" type dual-bucket-bandwidth
                                                granularity subscriber create
        description "Deflt downstream BW policer for 25Mbps Subs"
        rate 25000
    mbs
```

The AQP entry below will act as a default AQP policy since it does not include application or IP Header match criteria:

```
BNG>config>app-assure# group 1:1 policy
    app-qos-policy
        entry 500 create
            description "Deflt downstream BW policer for 25Mbps Subs"
            match
                traffic-direction network-to-subscriber
                characteristic "access-rate" eq "25M"
            exit
            action
                 bandwidth-policer "DefltPol-Sub-BW-DS-25Mbps"
            exit
            no shutdown
        exit
```

Note: A similar configuration can be implemented for the 100Mbps access rate service option.

Default Flow-Count-Limit Policing – 25Mbps AA-Sub

```
BNG>config>app-assure# group 1
    policer "DefltPol-Sub-FlowCount-US-25Mbps" type flow-count-limit
                                                granularity subscriber create
        description "Deflt policer to limit active upstream flows for 25Mbps Subs”
        flow-count 10000
        action permit-deny
    exit
    policer "DefltPol-Sub-FlowCount-DS-25Mbps" type flow-count-limit
                                                granularity subscriber create
        description "Deflt policer to limit active downstream flows for 25Mbps Subs”
        flow-count 10000
        action permit-deny
```

```
    exit
```

The AQP entry below will act as a default AQP policy since it does not include application or IP Header match criteria:

```
BNG>config>app-assure# group 1:1 policy app-qos-policy
    entry 510 create
        description " Deflt policer to limit active upstream flows for 25Mbps Subs"
        match
            traffic-direction subscriber-to-network
            characteristic "access-rate" eq "25M"
        exit
        action
            flow-count-limit "DefltPol-Sub-FlowCount-US-25Mbps"
        exit
        no shutdown
    exit
    entry 515 create
        description " Deflt policer to limit active downstream flows for 25Mbps Subs
"
        match
            traffic-direction network-to-subscriber
            characteristic "access-rate" eq "25M"
        exit
        action
            flow-count-limit "DefltPol-Sub-FlowCount-DS-25Mbps"
        exit
        no shutdown
    exit
```

Note: A similar configuration can be implemented for the 100Mbps access rate service option.

Default Flow-Rate-Limit Policing – 25Mbps AA-Sub

```
BNG>config>app-assure# group 1
    policer "DefltPol-Sub-FlowRate-US-25Mbps" type flow-rate-limit
                                                  granularity subscriber create
        description "Deflt policer to limit upstream flow setup rate for 25Mbps Subs
"
        rate 200
        action permit-deny
    exit
    policer "DefltPol-Sub-FlowRate-DS-25Mbps" type flow-rate-limit
                                                granularity subscriber create
        description "Deflt policer to limit downstr flow setup rate for 25Mbps Subs"
        rate 200
        action permit-deny
    exit
```

The AQP entry below will act as a default AQP policy since it does not include application or IP Header match criteria:

```
BNG>config>app-assure# group 1:1 policy app-qos-policy
    entry 520 create
        description "Deflt policer to limit upstream flow setup rate for 25Mbps Subs
"
        match
            traffic-direction subscriber-to-network
            characteristic "access-rate" eq "25M"
        exit
        action
            flow-rate-limit "DefltPol-Sub-FlowRate-US-25Mbps"
        exit
        no shutdown
    exit
    entry 525 create
        description "Deflt policer to limit downstr flow setup rate for 25Mbps Subs"
        match
            traffic-direction network-to-subscriber
            characteristic "access-rate" eq "25M"
        exit
        action
            flow-rate-limit "DefltPol-Sub-FlowRate-DS-25Mbps"
        exit
        no shutdown
    exit
```

Note: A similar configuration can be implemented for the 100Mbps access rate service option.

Application BW Policing (Per Subscriber)

The configuration example below provides a per AA subscriber peer-to-peer rate limit of 1Mbps. It does not include the app-profile configuration since the ASO characteristic and values can be either statically configured within the app-profile or dynamically signaled through RADIUS or Gx using ASO overrides.

AA subscribers with service characteristic "P2P-Sub-DL" value of "1M" will have a bandwidth policer of 1Mbps applied to peer to peer traffic in the network to subscriber direction:

```
BNG>config>app-assure# group 1
    policer "P2P-Sub-DL-1M" type single-bucket-
bandwidth granularity subscriber create
        description "Per-
subscr BW policer to limit P2P downstream traffic to 1Mbps"
        rate 1000
        mbs 19
```

```
        action permit-deny
    exit

BNG>config>app-assure# group 1:1 policy
    app-service-options
        characteristic "P2P-Sub-DL" create
            value "10M"
            value "1M"
            value "unlimited"
            default-value "unlimited"
        exit

BNG>config>app-assure# group 1:1 policy app-qos-policy
    entry 30 create
        description "Per-subscr BW policer to limit P2P downstream traffic to 1Mbps"
        match
            app-group eq "Peer to Peer"
            traffic-direction network-to-subscriber
            characteristic "P2P-Sub-DL" eq "1M"
        exit
        action
            bandwidth-policer "P2P-Sub-DL-1M"
        exit
        no shutdown
    exit
```

Conclusion

This example provides detailed information to properly configure and use app-profiles, ASOs and AQPs to successfully configure application policy control rules using Application Assurance.

Application Assurance — Asymmetry Removal

This chapter describes Application Assurance asymmetry removal configurations.

Topics in this chapter include:

Applicability

This chapter was originally written for and configured on release 11.0.R1. The CLI in the current edition corresponds to release 14.0.R4. The Application Assurance Redundancy Protocol (AARP) requires the FP2 or higher, in support for the spoke SDP divert feature.

The pre-requisites for this chapter are a base understanding of AA configuration and operation for single homed deployments. This chapter applies to dual-homed SAPs and spoke SDPs configurations, in a business or residential AA context. AARP is not used for ESM AA subscribers.

Overview

This chapter is intended for Application Assurance (AA) network architects and engineers. It provides best practices recommendations to configure AA Asymmetry Removal.

Asymmetry means that the two directions of a traffic flow (to-sub and from-sub) take different paths through the network. Asymmetry removal is a means of eliminating traffic asymmetry between a set of dual-homed SAP or spoke SDP endpoints. This can be across endpoints within a single node or across a pair of inter-chassis link connected routers, which is the topology explained in this chapter. Asymmetry removal ensures all packets of a dual-homed AA subscriber are diverted to an AA ISA in order to achieve accurate per subscriber traffic identification and policy enforcement.

Traffic asymmetry is created when there are dual-homed links for a service, and the links are simultaneously carrying traffic. Asymmetry removal for transit subscribers must be implemented in the first routed hop on the network side of the subscriber management point, so there will be a deterministic and fixed SAP/spoke SDP representing the downstream subscriber management node. This ensures there are no more than two paths that the flows can take, both covered by the asymmetry removal solution.

Configuration

Application Assurance Redundancy Protocol (AARP) provides the data plane connectivity for dynamically keeping a dual-homed AA subscriber's traffic on the same ISA-AA for AA processing. An AARP instance is configured between the dual-homed routers to establish connectivity with the same AARP instance number on each node.

When asymmetry exists between dual-chassis redundant systems, Ipipe spoke SDPs are used to interconnect these services between peer nodes over an Inter-Chassis Link (ICL). The following sections explain the configuration and operation of the services for use with the Application Assurance Redundancy Protocol.

AARP Service Configuration

The following services must be configured to establish communications between the AARP instances in each of the paired nodes.

- Network topology is a VPRN (or IES) service configured in each node, with a dual-homed SAP from each node to a downstream access element.
- Assumes starting point with AA ISAs installed with identical AA policy and divert enabled in each node.
- Also, the system needs basic routing and LDP configuration for the SDP and the spoke SDPs to be established.

Figure 12 **Application Assurance Asymmetry Removal Topology**

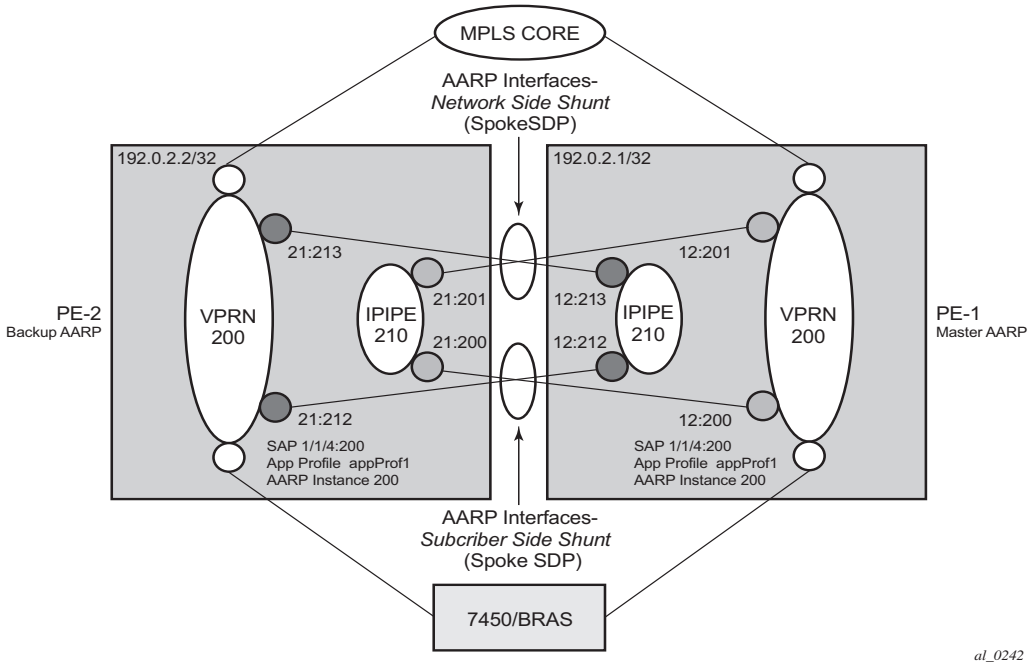

Table 4 **Application Assurance Asymmetry Removal Topology**

On PE-2	On PE-1
system ip: 192.0.2.2	system ip: 192.0.2.1
dual-homed service: 200	dual-homed service: 200
dual-homed sap: 1/1/4:200	dual-homed sap: 1/1/4:200
app-profile diverting: yes	app-profile diverting: yes

Configuration Commands for AARP

To enable AARP, AARP instances and AARP interfaces on both nodes must be configured. AARP operation has the following dependencies between the nodes:

- Shunt links configured and operationally up, both subscriber side shunt and network side shunt.
- Peer communications established between nodes, AARP instance operational status will be up when peers are communicating.

- Dual-homed sap/spoke SDP configured with a unique AARP instance (matched by dual-homed interface).
- App-profile configured against sap/spoke SDP with divert enabled (making the sub an aa-sub). The app-profile is the trigger to divert the traffic in the node with the active AARP instance to one of the ISAs in that node, per normal AA divert behavior.

Begin with PE-2:

```
configure
    application-assurance
        aarp 200 create
            description "aarp protecting a dual-homed sap"
            priority 100
            peer 192.0.2.1
            no shutdown
        exit
    exit
exit
```

Ipipe shunt configuration

```
configure
    service
        sdp 21 mpls create
            far-end 192.0.2.1
            ldp
            keep-alive
                shutdown
            exit
            no shutdown
        exit
        ipipe 210 customer 1 vc-switching create
            service-mtu 1556
            spoke-sdp 21:200 create
                aarp 200 type subscriber-side-shunt
                no shutdown
            exit
            spoke-sdp 21:201 create
                aarp 200 type network-side-shunt
                no shutdown
            exit
            no shutdown
        exit
    exit
exit
```

Dual-homed and Interface Shunt Configuration

```
configure
    service
        vprn 200 customer 1 create
            description "VPRN 200 Dual Homed Routed Service"
```

```
            aarp-interface "subside_1" create
                spoke-sdp 21:212 create
                    aarp 200 type subscriber-side-shunt
                    no shutdown
                exit
            exit
            aarp-interface "netside_1" create
                spoke-sdp 21:213 create
                    aarp 200 type network-side-shunt
                    no shutdown
                exit
            exit
            interface "int-BRAS-1" create
                sap 1/1/4:200 create
                    aarp 200 type dual-homed
                    app-profile "app-prof-1"
                exit
            exit
            no shutdown
        exit
    exit
exit
```

Then similarly configure the associated AARP configuration on **PE-1**:

```
configure
    application-assurance
        aarp 200 create
            description "aarp protecting a dual-homed sap"
            priority 200
            peer 192.0.2.2
            no shutdown
        exit
    exit
exit
```

Ipipe Shunt Configuration

```
configure
    service
        sdp 12 mpls create
            far-end 192.0.2.2
            ldp
            keep-alive
                shutdown
            exit
            no shutdown
        exit
        ipipe 210 customer 1 vc-switching create
            service-mtu 1556
            spoke-sdp 12:212 create
                aarp 200 type subscriber-side-shunt
                no shutdown
            exit
            spoke-sdp 12:213 create
                aarp 200 type network-side-shunt
                no shutdown
```

```
            exit
            no shutdown
        exit
    exit
exit
```

Dual-homed and Interface Shunt Configuration

```
configure
    service
        vprn 200 customer 1 create
            aarp-interface "subside_1" create
                spoke-sdp 12:200 create
                    aarp 200 type subscriber-side-shunt
                    no shutdown
                exit
            exit
            aarp-interface "netside_1" create
                spoke-sdp 12:201 create
                    aarp 200 type network-side-shunt
                    no shutdown
                exit
            exit
            interface "int-BRAS-1" create
                sap 1/1/4:200 create
                    description "AA enabled SAP"
                    aarp 200 type dual-homed
                    app-profile "app-prof-1"
                exit
            exit
            no shutdown
        exit
    exit
exit
```

Show Commands for AARP

Verify correct configuration on each node. The following output displays the example configuration for PE-1.

Starting with the AARP instance in each node, verify that the AARP instance operational state is up (if everything is properly configured as intended):

```
*A:PE-1# show application-assurance aarp 200
===============================================================================
AARP Instance 200
===============================================================================
Description    : aarp protecting a dual-homed sap
Admin State    : Up                     Oper State     : Up

Local IP       : 192.0.2.1              Peer IP        : 192.0.2.2
Local State    : master                 Peer State     : backup
Local Priority : 200                    Peer Priority  : 100
```

```
Local Flags    : none
Peer Flags     : none
Peer End-Point : none

Master Selection Mode       : minimizeSwitchovers

-------------------------------------------------------------------------------
Service References
-------------------------------------------------------------------------------
Service              Reference            Reference Type
-------------------------------------------------------------------------------
VPRN 200             1/1/4:200            Dual-Homed
Ipipe 210            12:212               Subscriber-Side Pipe Shunt
Ipipe 210            12:213               Network-Side Pipe Shunt
VPRN 200             12:200               Subscriber-Side AARP-Interface Shunt
VPRN 200             12:201               Network-Side AARP-Interface Shunt
-------------------------------------------------------------------------------
No. of service references: 5
-------------------------------------------------------------------------------
===============================================================================
*A:PE-1#
```

Verifying that the AARP instance is up is an indication that the dual-node communications for AARP is working (instance, shunts, etc.). In addition, in the preceding output, verify on both PE nodes that the intended SAPs are dual-homed for that instance.

Now a detailed review of the configured AARP shunt infrastructure services can be shown to make sure they are all properly configured with intended AARP parameters (such as AARP ID and Type on the network and subscriber side shunts) as displayed in the following output:

```
*A:PE-1# show service id 210 all

===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 210                Vpn Id            : 0
Service Type      : Ipipe
Name              : (Not Specified)
Description       : (Not Specified)
Customer Id       : 1                  Creation Origin   : manual
Last Status Change: 10/03/2016 11:45:51
Last Mgmt Change  : 10/03/2016 11:45:51
Admin State       : Up                 Oper State        : Up
MTU               : 1556
Vc Switching      : True
SAP Count         : 0                  SDP Bind Count    : 2
CE IPv4 Discovery : n/a                Keep address      : No
CE IPv6 Discovery : n/a                Stack Cap Sig     : n/a

Eth Legacy Fault Notification
-------------------------------------------------------------------------------
Recovery Timer    : 10.0 secs          Admin State       : outOfService
```

```
-------------------------------------------------------------------------------
ETH-CFM service specifics
-------------------------------------------------------------------------------
Tunnel Faults     : ignore


-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 12:212  -(192.0.2.2)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 12:212                  Type              : Spoke
Spoke Descr     : (Not Specified)
Split Horiz Grp   : (Not Specified)
VC Type           : Ipipe                   VC Tag            : 0
Admin Path MTU    : 0                       Oper Path MTU     : 1556
Delivery          : MPLS
Far End           : 192.0.2.2
Tunnel Far End    : 192.0.2.2               LSP Types         : LDP
Hash Label        : Disabled                Hash Lbl Sig Cap  : Disabled
Oper Hash Label   : Disabled
Entropy Label     : Disabled

Admin State       : Up                      Oper State        : Up
MinReqd SdpOperMTU : 1556
Acct. Pol         : None                    Collect Stats     : Disabled
Ingress Label     : 262141                  Egress Label      : 262139

--- snipped ---

Application Profile: None
Transit Policy    : None
AARP Id           : 200
AARP Type         : subscriber-side-shunt

--- snipped ---

-------------------------------------------------------------------------------
IPIPE Service Destination Point specifics
-------------------------------------------------------------------------------
Configured CE IPv4 Addr: n/a                Peer CE IPv4 Addr : 0.0.0.0

-------------------------------------------------------------------------------
 Sdp Id 12:213  -(192.0.2.2)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 12:213                  Type              : Spoke
Spoke Descr     : (Not Specified)
Split Horiz Grp   : (Not Specified)
VC Type           : Ipipe                   VC Tag            : 0
Admin Path MTU    : 0                       Oper Path MTU     : 1556
Delivery          : MPLS
Far End           : 192.0.2.2
Tunnel Far End    : 192.0.2.2               LSP Types         : LDP
Hash Label        : Disabled                Hash Lbl Sig Cap  : Disabled
Oper Hash Label   : Disabled
Entropy Label     : Disabled
```

```
Admin State        : Up                       Oper State        : Up
MinReqd SdpOperMTU : 1556
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : 262140                   Egress Label      : 262138

--- snipped ---

Application Profile: None
Transit Policy     : None
AARP Id            : 200
AARP Type          : network-side-shunt

--- snipped ---

===============================================================================
*A:PE-1#
```

Next, the configuration of the VPRN service of the dual-homed SAP can be reviewed to ensure it reflects the attached endpoints for the shunt Ipipe spoke SDPs:

```
*A:PE-1# show service id 200 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 200                  Vpn Id            : 0
Service Type      : VPRN
Name              : (Not Specified)
Description       : (Not Specified)
Customer Id       : 1                    Creation Origin   : manual
Last Status Change: 10/03/2016 11:45:51
Last Mgmt Change  : 10/03/2016 11:45:51
Admin State       : Up                   Oper State        : Up

Route Dist.       : 64496:200            VPRN Type         : regular
Oper Route Dist   : 64496:200
Oper RD Type      : configured
AS Number         : None                 Router Id         : 192.0.2.1
ECMP              : Enabled              ECMP Max Routes   : 1
Max IPv4 Routes   : No Limit
Auto Bind Tunnel
Resolution        : disabled
Max IPv6 Routes   : No Limit
Ignore NH Metric  : Disabled
Hash Label        : Disabled
Entropy Label     : Disabled
Vrf Target        : target:64496:200
Vrf Import        : None
Vrf Export        : None
MVPN Vrf Target   : None
MVPN Vrf Import   : None
MVPN Vrf Export   : None
Car. Sup C-VPN    : Disabled
Label mode        : vrf
BGP VPN Backup    : Disabled
BGP Export Inactv : Disabled

SAP Count         : 1                    SDP Bind Count    : 2
```

```
VSD Domain        : <none>

--- snipped ---

-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
 Sdp Id 12:200  -(192.0.2.2)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 12:200                  Type              : Spoke
Spoke Descr     : (Not Specified)
VC Type           : n/a                     VC Tag            : n/a
Admin Path MTU    : 0                       Oper Path MTU     : 1556
Delivery          : MPLS
Far End           : 192.0.2.2
Tunnel Far End    : 192.0.2.2               LSP Types         : LDP
Hash Label        : Disabled                Hash Lbl Sig Cap  : Disabled
Oper Hash Label   : Disabled
Entropy Label     : Disabled

Admin State       : Up                      Oper State        : Up

--- snipped ---

Application Profile: None
Transit Policy    : None
AARP Id           : 200
AARP Type         : subscriber-side-shunt

--- snipped ---

-------------------------------------------------------------------------------
IPIPE Service Destination Point specifics
-------------------------------------------------------------------------------
Configured CE IPv4 Addr: n/a                Peer CE IPv4 Addr : 0.0.0.0

-------------------------------------------------------------------------------
 Sdp Id 12:201  -(192.0.2.2)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 12:201                  Type              : Spoke
Spoke Descr     : (Not Specified)
VC Type           : n/a                     VC Tag            : n/a
Admin Path MTU    : 0                       Oper Path MTU     : 1556
Delivery          : MPLS
Far End           : 192.0.2.2
Tunnel Far End    : 192.0.2.2               LSP Types         : LDP
Hash Label        : Disabled                Hash Lbl Sig Cap  : Disabled
Oper Hash Label   : Disabled
Entropy Label     : Disabled

Admin State       : Up                      Oper State        : Up

--- snipped ---

Application Profile: None
Transit Policy    : None
```

```
AARP Id            : 200
AARP Type          : network-side-shunt

--- snipped ---

*A:PE-1#
```

Continuing deeper into the same VPRN service show output, or using the following show command, it can be verified that the dual-homed SAP itself is properly configured and associated with that service and AARP instance:

```
*A:PE-1# show service id 200 sap 1/1/4:200 detail

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 200
SAP                : 1/1/4:200                Encap             : q-tag
Description        : AA enabled SAP
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 10/03/2016 11:45:51
Last Mgmt Change   : 10/03/2016 11:45:51
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

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

Application Profile: app-prof-1
Transit Policy     : None
AARP Id            : 200
AARP Type          : dual-homed

Oper Group         : (none)                   Monitor Oper Grp  : (none)
Host Lockout Plcy  : n/a
Lag Link Map Prof  : (none)
Bandwidth          : Not-Applicable

--- snipped ---

===============================================================================
*A:PE-1#
```

Network to Subscriber Traffic Flow

When the AARP is operationally up, AARP tracks which ISA is the master ISA for each dual-homed AARP instance and uses the inter-chassis services (spoke SDP AARP shunts) to move all traffic for each instance traffic to the node with the Master ISA.

Looking at traffic in the network to subscriber direction (Figure 13):

- Traffic arriving on PE-1 is diverted to the local master ISA, processed, then proceeds to the egress SAP.
- Traffic arriving on PE-2 with the backup AARP interface is sent to the master node for AA processing. The ingress FP forwards packets to network-side-shunt AARP interface for remote AA divert.
- Arriving on PE-1, the packets on the AARP Ipipe are diverted to the master ISA where the packets are processed as if this traffic was traveling in the to-sub direction towards the dual-homed endpoint on PE-1, then returned to PE-2.
- Entering PE-2, the traffic from the subscriber side shunt interface is not diverted to ISAs in that node and egresses on the AARP instance SAP.

With this behavior, traffic always returns to the original ingress node before egressing toward the subscriber (network path for the flows are not modified).

Figure 13 **Network to Subscriber Traffic Flow**

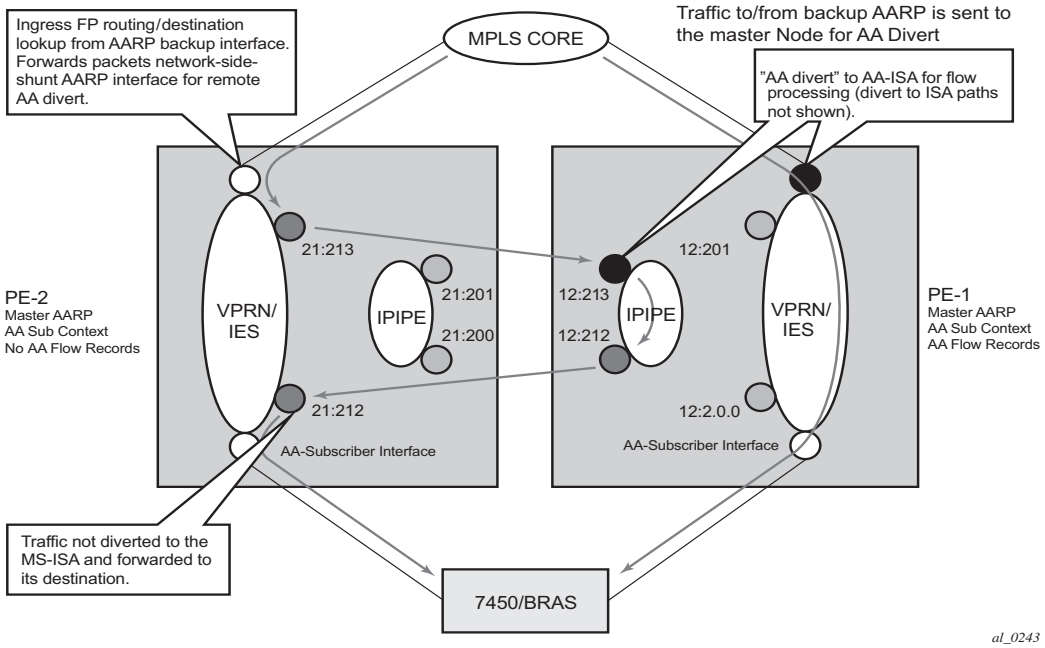

Subscriber to Network Traffic Flow

Looking at traffic in the subscriber to network direction (Figure 14):

- Traffic arriving on PE-1 is diverted to the local master ISA, processed, then proceeds to the egress SAP.
- Traffic arriving on PE-2 with the backup AARP ISA is sent to the master node for AA processing (not diverted to an ISA in PE-2). The ingress FP forwards packets to subscriber-side-shunt AARP interface for remote AA divert.
- Arriving on PE-1, the packets on the AARP Ipipe are diverted to the master ISA where the packets are processed as if the traffic was flowing in the from-sub direction on the dual-homed endpoint, then returned to PE-2 over the Ipipe's AARP subscriber-side-shunt.
- Entering PE-2, the traffic from the network side shunt interface is forwarded by the IES/VPRN service to its destination.

Figure 14 **Subscriber to Network Traffic Flow**

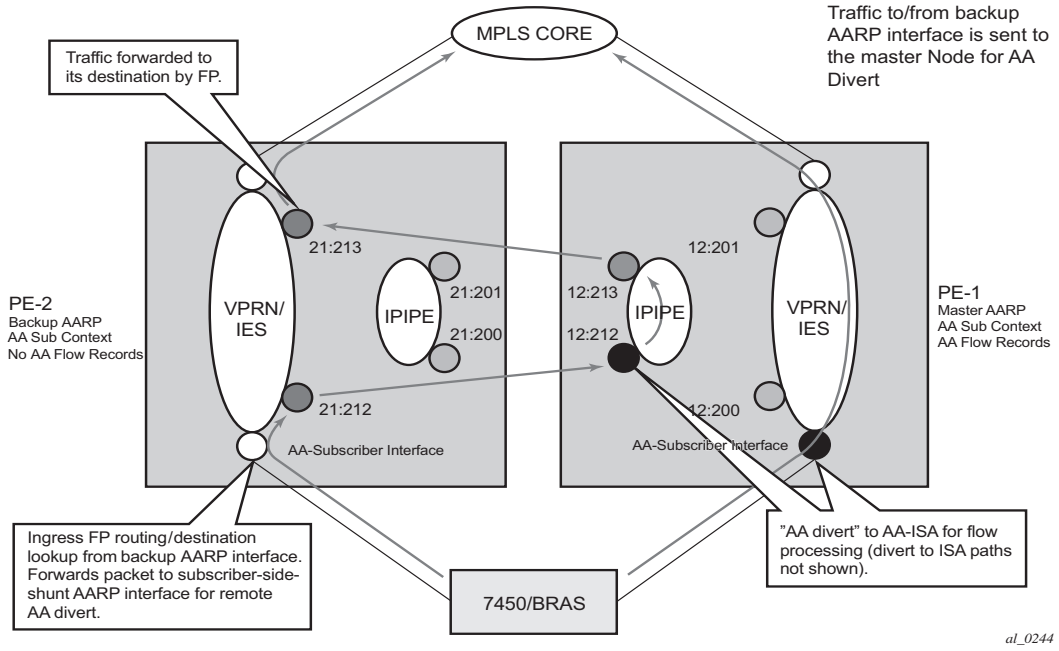

Typical Configuration Mistakes

Operators configuring AARP can make some typical mistakes listed below that will keep the AARP instance in Operational State down:

- The spoke SDP AARP shunt instances' IDs must be aligned with the respective spoke SDP on the peer node: if not, it will result in a flag indicating **shunt mismatch** in the show output.
- Ipipe service MTU alignment — The Ipipe service MTU values must be the same in both nodes, otherwise it will result in the services be in operational status UP, but the AARP instance will remain down.

Conclusion

This chapter is intended for Application Assurance (AA) network architects and engineers to provide the information required to understand and configure dual-node asymmetry removal following the intended service configuration as used by the AARP implementation.

Application Assurance — Best Practices for ISA and Host IOM Overload Protection

This chapter provides information about Application Assurance best practices for ISA and host IOM overload protection.

Topics in this chapter include:

Applicability

This configuration note is applicable to all 7750 SR/SR-c and 7450 ESS chassis supporting Application Assurance (AA).

The configuration was tested on release 12.0.R4.

Overview

The Multi-Service Integrated Services Adapter (MS-ISA) is a processing resource module installed in a7x50 SR/ESS system on an ISA host IOM. This example describes the best practices for configuration and monitoring of the system to ensure proper engineering of the system resources involved in AA ISA capacity planning.

As shown in Figure 15, traffic is diverted to an AA ISA by provisioning of an application profile (app-profile) for a subscriber or SAP service context. SR OS then automatically handles traffic diversion for both directions of traffic for that AA subscriber context, through one of the AA ISAs in the AA group where that app-profile is defined.

Figure 15 **System Packet Datapath to AA ISA**

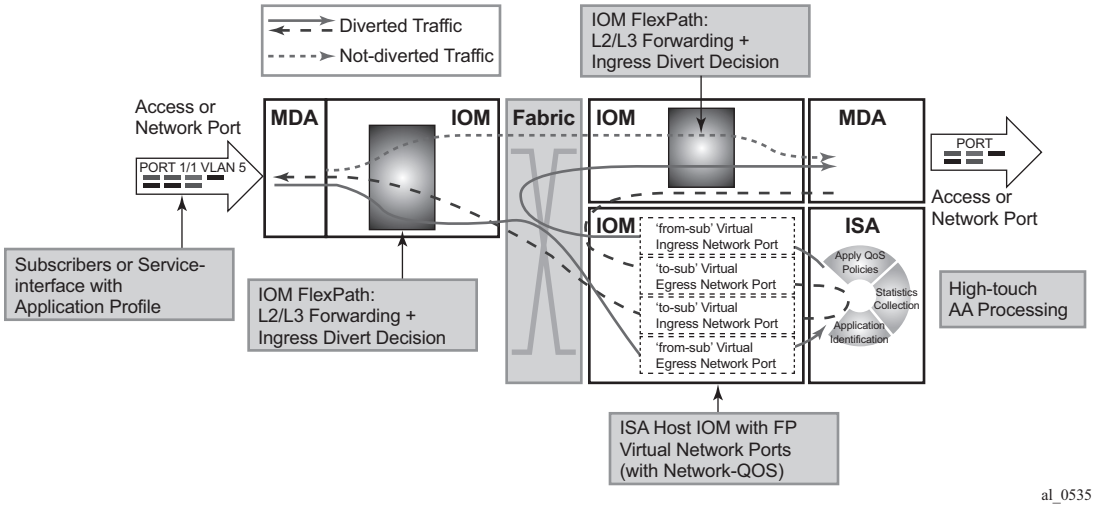

The following elements in the 7x50 SR/ESS system must be properly engineered for any given AA deployment. Each element is described in this section:

1. ISA capacity cost and load balancing across ISAs.
2. ISA host IOM network egress QoS. Host IOM egress network ports weighted-average shared buffer pool thresholds (within the egress QoS configuration for each group) are used for overload cut-through processing.
3. ISA resources and statistics collection.
   - Flows
   - Traffic volume (bandwidth)
   - Subscribers
   - Flow setup rate
   - ISA overload cut-through
   - ISA default subscriber policies

ISA Capacity Planning Approach

This example illustrates an approach to the configuration of the 7x50 SR/ESS AA system to address these considerations:

- IOM/ISA-AA network egress QoS configuration should be designed to treat the ISA as a network port with normal network port maximum delay (by MBS).

- Within the ISA, fair access to the ISA-AA bandwidth and flow resources must be ensured: it is recommended that default Application QoS Policy (AQP) policy entries be configured limiting bandwidth and flow resources per AA subscriber.
- Thresholds for SNMP alerts that indicate a high load on ISA processing should be configured: capacity cost, flow, bandwidth.
- Capacity tracking in live deployments should be performed for parameters that can affect overload: flow setup rate, bandwidth, and subscriber-count per ISA.
- Use of other scale related consumable AA resources against system maximum limits. This includes parameters such as statistics records, transit-ip table entries, and transit-prefix TCAM entries, which should be planned and periodically tracked. These limits will not affect overload of the ISA but may affect intended service operation.
- For recommendations of the specific parameters to watch in a given deployment as well as the values of the system limits for a given release, contact your regional support organization.

AA Overload and Resource Monitoring

Overload is a condition where the total packet processing requirements for traffic arriving on a given ISA exceeds the available resources, resulting in the host IOM egress buffers reaching a configured “overload” threshold. Above this threshold, the ISA can be configured to forward excess traffic (called overload cut-through). If cut-through is not enabled and the overload condition continues, the egress queue MBS threshold will eventually be reached, after which packets will be discarded. Even if overload cut-through is enabled, any egress traffic that exceeds the maximum bus capacity of the ISA queue discard threshold will be discarded.

ISA capacity overload events are supported within the system resource monitoring and logging capabilities if the traffic and resource load crosses any of the following high and low load thresholds on a per-ISA basis. Exceeding one of these thresholds does not in itself indicate an overload state.

a. Host IOM egress network ports weighted-average shared buffer pool thresholds (within the egress QoS configuration for each AA group) are used for triggering and removing overload cut-through processing. Care should be taken in the configuration of these buffers, as the IOM flexpath has significant buffer capacity that can result in latency larger than the network SLA acceptable guidelines. A properly engineered configuration will have large enough buffering to not trigger ISA overload unnecessarily (due to normal bursts with a reasonable traffic load) but will not incur excessive latency prior to triggering the overload state.

b. ISA capacity cost: if the capacity cost of all subscribers on the ISA exceeds a threshold, an event is raised but the overload condition is not set (unless other resources are exhausted). ISA overload or traffic cut-through does not occur simply if capacity cost is exceeded. It is used to capacity plan an intended load for the ISA, proportional to resource use per subscriber, in order to generate events prior to overload to allow appropriate action to mitigate the resource consumption (such as provisioning more ISAs).

c. Flow table consumption (number of allocated flow resources in use): the flow table high-watermark threshold warnings are for proactive notification of a high load. The ISA will cut-through new flows when the “flow resources in use” is at the maximum flow limit. Reaching the flow limit does not generate backpressure to the IOM, nor is the ISA considered in an overload state. Flow usage thresholds are different from bit-rate/packet-rate/flow-setup-rate thresholds in that when the flow table high-watermark threshold is exceeded, the ISA will no longer be operating as application-aware for the flows with no context. The default subscriber policy is applied to traffic that required a flow record but was unable to allocate one, which is a similar behavior to overload cut-through.

   The following terms are used to describe flow resources:

   – Maximum flows: the maximum AA flow table size for a given release.
   – Flows: on the show screens, the “flows” field is an indication of the number of unique 5-tuple entries in the flow table. This includes active and inactive flows; inactive will age out of the table after a period of inactivity that is dependent on the protocol used.
   – Active flows: the number of flows with traffic in the current reporting interval.
   – Flow resources in use: the number of allocated flows in the flow table. This number is greater than the number of active flows, reflecting inactive flows and flows pre-allocated for some dynamic protocols (control + data channels) and for some UDP traffic.

d. Traffic volume: traffic rate in bytes/sec and packets/sec is the dominant cause of ISA overload in most network scenarios, when the ISA is presented with more traffic than it can process. This results in the ISA internal ingress buffers reaching a threshold that causes backpressure to the IOM egress queues (toward the ISA), allowing the ISA to process the packets it already has. This internal backpressure mechanism is normal behavior, allowing burst tolerance at the IOM-to-ISA interface; thus backpressure is not in itself an indication of overload. Overload occurs when the bursts or the load of traffic is sustained long enough to reach the ISA host IOM network port egress weighted-average shared buffer threshold. The actual amount of traffic that can be passed through an ISA is dependent on the application traffic mix, flow density, and AA policy configurations and will vary by network type and by region. The bit-rate and packet-rate watermarks can be used to provide event notification when the traffic rates exceed planning expectations.

e. Flow setup rate: this is generally proportional to total traffic volume, and as such can be a factor in ISA overload. The flow setup rate is the rate at which new flows are presented to the ISA, each resulting in additional tasks that are specific to flow state creation; thus the ISA has a sensitivity to flow setup rates as fewer cycles are available for other datapath tasks when the flow setup rate is high. In residential networks, flow setup rates of 3 k to 5 k flows/sec per Gbps of traffic are common. The flow setup watermarks can be used to provide event notification when the rate exceeds planning expectations.

ISA Overload Models

For an ISA overload strategy, there are two design options for configuring the overload behavior of the system:

- Host IOM egress discards: in this model, the philosophy is to treat AA packet processing resources in the same way as a network interface (of somewhat variable capacity depending on the traffic characteristics). When too much traffic is presented to the ISA, it backpressures the host IOM egress, which will buffer packets. If the egress buffer thresholds are exceeded, the ISA will discard according to the egress QoS slope policy. This is configured by not enabling **isa-overload-cut-through** and use of appropriate egress QoS policies. Firewall or session filter deployments may use this model.
- Overload cut-through: the ISA group can be enabled to cut-through some traffic if an overload event occurs, triggered when the IOM network port weighted-average queues depth exceeds the weighted-average shared high-watermark threshold. In this ISA state, some packets are cut-through from application analysis but retain subscriber context with the default subscriber policy applied. This mode of deployment is intended for situations where it is preferable to forward packets even if not identified by AA than to drop/discard the packet. For example, if AA is providing value-added services (VAS) such as In-Browser Notification (IBN), analytics, or traffic rate limiting, this would usually be the preferred model as the underlying service should be preserved even if capacity to provide the VAS is not available.

Note that even with overload cut-through enabled, there is a hardware-based maximum ISA throughput of approximately 11 Gbps for MS-ISA and 40 Gbps for MS-ISA2. If this is exceeded on a sustained basis, IOM egress discards may still occur.

Understanding Packet and Protocol Cut-Through

Traffic can be cut-through the ISA-AA card on a packet-by-packet basis, in which case packets do not go through AA identification and subscriber application policy. The conditions that trigger cut-through include:

- Overload (IOM egress network port weighted-average shared buffer threshold): excess traffic bypasses all AA processing except for the default subscriber policy
- Non-conformant IP packet: traffic bypasses all AA processing except IP protocol checks and the default subscriber policy. Optionally, these packets can be discarded in AA.
- Flow table full: for new 5-tuples sent to the ISA, if the flow table is full, the packets are cut-through the ISA and only the default subscriber policy is applied.

Note: The default subscriber policy is a set of AQP rules that apply AQP match criteria limited to Application Service Options (ASO), aa-sub, and traffic-direction starting with the first packet of a flow, with no match conditions based on AA identification (application, app-group, charging-group, IP header). Packets will be either denied_by_default_policy or cut_through_by_default_policy, depending on the policer action configuration in the AQP rules.

For cut-through traffic, no flow records exist but it is counted under per-subscriber protocol statistics as one of the following counters, depending on the case:

- cut_through — Statistics for any packet that could not map to a flow, but that has a valid subscriber ID. This can be an error packet, fragmented out-of-order, no flow resource, invalid TCP flags, etc. This is the most important count for indicating overload cut-through, as it counts all traffic in overload cut-through mode (when the weighted-average threshold has been crossed).
- denied_by_default_policy — Packets that are dropped due to a default policy with a flow-based policer (flow rate or flow count) with action discard.
- cut_through_by_default_policy — Packets that failed to pass flow-based policers with an action of priority-mark.

An example of overload cut-through statistics in the CLI is shown below:

```
A:BNG# show application-assurance group 1 protocol count
===============================================================================
Application-Assurance Protocol Statistics
===============================================================================
Protocol                          Disc          Octets       Packets      Flows
-------------------------------------------------------------------------------
advanced_direct_connect             0%               0             0          0
aim                                 0%               0             0          0
```

amazon_video	0%	0	0	0
ares	0%	0	0	0
bbm	0%	0	0	0
betamax_voip	0%	0	0	0
bgp	0%	0	0	0
bittorrent	0%	678428534	5322929	1036129
cccam	0%	0	0	0
citrix_ica	0%	0	0	0
citrix_ima	0%	0	0	0
cnnlive	0%	0	0	0
cups	0%	0	0	0
cut_through	0%	5299435739	10603771	0
cut_through_by_default_policy	0%	0	0	0
cvs	0%	0	0	0
daap	0%	0	0	0
dcerpc	0%	0	0	0
denied_by_default_policy	0%	0	0	0

Configuration

This example illustrates a typical, configuration of a 7x50 system for AA each of the configuration topics.

AA Traffic Load Test Environment

Application assurance identifies every byte and every packet of hundreds of real-world applications using per-flow stateful analysis techniques. It is a challenge to find test equipment that can accurately emulate full scale (10 Gbps to 40 Gbps) with traffic mixes and flow behaviors representing hundreds of thousands of end users with application clients across a range of devices. Some specialized stateful test equipment can emulate large traffic rates, but even the best will have equipment-specific patterns and behaviors not representative of live traffic. Therefore, the best scenario to engineer the AA overload configuration is by iteration in live deployments: setting an initial target and modifying the configuration based on ISA performance under load.

For a lab test of ISA throughput and loading, Nokia uses stateful test equipment which supports emulation of various service provider traffic mix profiles suitable for generating overload conditions; however, it is outside the scope of this document to configure AA throughput tests.

The operator should be aware that use of unrealistic, non-stateful traffic generators can result in a high level of unknown traffic, with the ISA performance impacted by continually trying to identify large numbers of packets of no real application type. This, combined with cut-through for invalid IP packets, can result in ISA overload and traffic cut-through (due to overload or invalid IP packets) at traffic levels not representative of actual ISA performance on real traffic.

ISA Capacity Cost and Load Balancing Across ISAs

These AA group-level commands define the load balancing parameters within an ISA group.

```
*A:BNG# configure isa
    application-assurance-group 1 aa-sub-scale residential create
        no description
        no fail-to-open
        isa-capacity-cost-high-threshold 304000
        isa-capacity-cost-low-threshold 272000
        partitions
        divert-fc be
        no shutdown
    exit
```

The following should be noted related to this configuration:

- Up to 7 primary and 1 backup ISAs are allowed. If the AA services are considered “value added” and not part of a paid service, backups are usually not used since the “fail to fabric” capability keeps the underlying service running.
- The default behavior in case of ISA failure is “no fail-to-open”, which means “fail-to-wire”; if an ISA fails, traffic is forwarded as if **no divert** was configured
- Threshold for sending capacity-cost SNMP traps: the unit used for capacity cost is a variable defined in the network design; in this example, it is expressed in Mbps of the subscriber total BW UP+DOWN with a high watermark set to 7600 Mbps × 40 = 304000 (where 40 is an oversubscription ratio). The low watermark is equal to 6800 Mbps × 40 = 272000.
- Partitions should always be enabled to configure additional policies in the future (for example, wifi/business)
- **divert-fc** configuration applies to the AA group: in this example, FC BE Internet is the only diverted FC; this is typical for AA residential and WLan-GW deployments. For VPN services, typically all datapath FCs are diverted to AA.

ISA-AA Host IOM - Network Egress Shared Memory and QoS

The amount of shared memory allocated per port, along with the network port egress QoS policy, determine the maximum delay for traffic diverted to Application Assurance.

This maximum network port delay is typically determined by the operator and must be used to define the proper QoS configuration to apply to the ISA-AA ports; this QoS configuration may be the same (typically) as what is applied to regular network ports on the 7x50 SR/ESS.

For IOM3-XP and FP2 and higher based line cards there is shared network egress memory per ISA-AA port, with the ISA-AA is represented by two network ports on the host IOM:

- “from-sub”: for traffic sent from the subscriber to the network
- “to-sub”: for traffic sent from the network to the subscriber

```
configure isa application-assurance-group 1
            qos
                egress
                    from-subscriber
                        pool
                            slope-policy "default"
                            resv-cbs default
                        exit
                        queue-policy "network-facing-egress"
                        port-scheduler-policy "network-facing"
                    exit
                    to-subscriber
                        pool
                            slope-policy "default"
                            resv-cbs default
                        exit
                        queue-policy "network-facing-egress"
                        port-scheduler-policy "network-facing"
                    exit
            exit
            no shutdown
```

For IOM3-XP and FP2 and higher based line cards the amount of shared memory reserved for each egress network port is determined by the speed of the port (10 Gbps for MS-ISA and 40 Gbps for MS-ISA2) and the **egr-percentage-of-rate** ratio configuration.

MS-ISA uses by default 1000% and 500% of the rate respectively for to-sub and from-sub ports, while MS-ISA2 uses by default 100% for both to-sub and from-sub ports.

It is typically recommended that these values be adjusted when MS-ISA and a high-speed Ethernet MDA are mixed on the same IOM3, since in this context the amount of shared memory allocated to the Ethernet MDA should be increased by reducing the MS-ISA network ports memory allocation ratio. If two MS-ISAs are installed on the same IOM3, the system will by default allocate 50% of the network egress shared memory to each ISA. In addition, an operator may adjust these values in case the actual network-to-subscriber versus subscriber-to-network ratio is significantly different in the production network, in order to achieve the expected maximum tolerated network delay.

The operator can modify the **egr-percentage-of-rate** per port using the command below:

```
A:BNG# configure port 1/2/fm-sub
A:BNG>config>port# info detail
----------------------------------------------
        modify-buffer-allocation-rate
            egr-percentage-of-rate 500
        exit
----------------------------------------------
A:BNG# configure port 1/2/to-sub
A:BNG>config>port# info detail
----------------------------------------------
        modify-buffer-allocation-rate
            egr-percentage-of-rate 1000
        exit
```

Network egress scheduling/queuing priority is for all ISAs within a group defined at the AA ISA group level

An example below with ISA-AA and 2 x 10G Eth MDA:

```
7750# configure port <slot>/<isa-aa-mda>/fm-sub
    modify-buffer-allocation-rate
        egr-percentage-of-rate 65

7750# configure port <slot>/<isa-aa-mda>/to-sub
    modify-buffer-allocation-rate
        egr-percentage-of-rate 130
```

In this example, the configuration defines:

- from-sub — Approximately 190 msec worth of buffer at 2500 Mbps.
- to-sub— Approximately 190 msec worth of buffer at 5000 Mbps.
- The buffer can be further refined from the network QoS policy.

For MS-ISA2, each MS-ISM flexpath will default the buffer allocation rate to 100%, which is a suitable value assuming that both modules in a slot are MS-ISA2 (which is the MS-ISM configuration), or that the I/O module has a similar traffic rate as the MS-ISA2 (which is also the case in the 10x10GE and 1x100GE versions of the MS-ISA2 line cards).

Configuring ISA Resources and Stats Collection

The following are the key consumable resources in an AA ISA:

- Flows
- Bandwidth
- Subscribers
- Flow setup rate

The AA group should be configured with watermark thresholds where each ISA will generate SNMP events when resources reach this level.

- Per-ISA-card resource usage watermarks trigger SNMP traps to the management system (5620 SAM)
- The values defined below can be refined based on the network characteristics in term of flows and bandwidth per ISA after the initial deployment

```
7750# configure application-assurance
-----------------------------------------------
        flow-table-low-wmark 90
        flow-table-high-wmark 95
        flow-setup-high-wmark 66500
        flow-setup-low-wmark 63000
        bit-rate-high-wmark 7600
        bit-rate-low-wmark 6800
```

In this example, the usage SNMP watermarks are configured for:

- Flow table: 95%/90% (maximum 4M flows on MS-ISA)
- Flow setup rate: configured to 95%/90% (of maximum 70k fps on MS-ISA)
- Bit rate/total diverted throughput

The **show>app-assure>group>status detail** command is used to display basic ISA health status:

- # aa-sub, active aa-sub, bitrate, flows in use, flow setup rate
- statistics for all ISAs combined or per ISA

```
A:BNG# show application-assurance group 1 status detail
  ===============================================================================
  Application-Assurance Status
  ===============================================================================
  Last time change affecting status : 05/30/2014 17:18:34
  Number of Active ISAs             : 4
  Flows                             : 214007945881
  Flow Resources In Use             : 2955164
  AA Subs Created                   : 70567
  AA Subs Deleted                   : 10544
  AA Subs Modified                  : 0
  Seen IP Requests Sent             : 0
  Seen IP Requests Dropped          : 0
  -------------------------------------------------------------------------------
                                      Current    Average    Peak
  -------------------------------------------------------------------------------
  Active Flows                      : 2911508    2769454    4582522
  Flow Setup Rate (per second)      : 33923      29400      67865
  Traffic Rate (Mbps)               : 7620       7238       22628
  Packet Rate (per second)          : 1254138    1182571    3044376
  AA-Subs Downloaded                : 69887      66129      70567
  Active Subs                       : 23131      19737      38114
  -------------------------------------------------------------------------------
                                      Packets              Octets
  -------------------------------------------------------------------------------
  Diverted traffic                  : 7437950197613        5530634242355947
  Diverted discards                 : 0                    0
      Congestion                    : 0                    0
      Errors                        : 0                    N/A
  Entered ISA-AAs                   : 7437950180191        5530634229794634
  Buffered in ISA-AAs               : 22                   29849
  Discarded in ISA-AAs              : 97790                47801217
      Policy                        : 0                    0
      Congestion                    : 0                    0
      Errors                        : 97790                47801217
  Modified in ISA-AAs
      Packet size increased         : 0                    0
      Packet size decreased         : 0                    0
  Errors (policy bypass)            : 28283549             21160338635
  Exited ISA-AAs                    : 7437950082379        5530634181963568
  Returned discards                 : 0                    0
      Congestion                    : 0                    0
      Errors                        : 0                    N/A
  Returned traffic                  : 7437950054070        5530634162337570
  ===============================================================================
```

This can also be run on a per-ISA basis:

```
show application-assurance group 1 status isa <slot/
port> detail
```

Note that for MS-ISA2, there is a maximum AA packet rate of 7 M pps; under most known traffic mix scenarios, the ISA should be safely below this packet rate when at maximum bandwidth throughput. However, it is worth periodically checking this value, because if the maximum packet rate is exceeded, and overload cut-through will result. (For MS-ISA, the maximum packet rate supported is high enough to not be feasible with realistic application-based traffic mixes).

The ISA aa-performance record should always be enabled in a network for capacity planning purposes in order to properly plan when to add new ISA cards if required and to monitor the network health:

```
*A:BNG>config>isa# info
----------------------------------------------
        application-assurance-group 1 aa-sub-scale residential create
            no description
            primary <slot/port>
            backup <slot/port>
            no fail-to-open
            isa-capacity-cost-high-threshold 304000
            isa-capacity-cost-low-threshold 272000
            partitions
            statistics
                performance
                    accounting-policy 7
                    collect-stats
                exit
            exit
            divert-fc be
            no shutdown
        exit
```

The commands highlighted in bold above will export information on the total traffic load and resource utilization of the ISA card:

- Flows — active flows, setup rates, resource allocation
- Traffic rates — bandwidth, packets
- Subscribers — active, configured, statistics resource allocation in use

The AA statistics collection configuration refers to accounting policies that are also defined in the 7x50 SR/ESS system:

```
*A:BNG>config# log
    file-id 7
        description "ISA Performance Stats"
        location cf2:
        rollover 15 retention 12
    exit
    accounting-policy 7
        description "ISA Performance Stats"
        collection-interval 15
        record aa-performance
```

```
            to file 7
            no shutdown
        exit
```

From the AA performance record the following fields in Table 5 can be used as to tracking ISA load in the reporting interval (typically a 15 to 60 minute period):

Table 5 **Tracking ISA Load in the Reporting Interval**

Record Name	Type	Description	Load planning use
dco	cumulative	octets discarded due to congestion in MDA	Should be 0; ISA internal congestion
dcp	cumulative	packets discarded due to congestion in MDA	Should be 0; ISA internal congestion
dpo	cumulative	octets discarded due to policy in MDA	Not related to load planning
dpp	cumulative	packets discarded due to policy in MDA	Not related to load planning
pbo	cumulative	octets policy bypass	Not used. Traffic was for an invalid subscriber and the group was "no fail-to-open"
pbp	cumulative	packets policy bypass	Not used. Traffic was for an invalid subscriber and the group was "no fail-to-open"
nfl	cumulative	number of flows	informative
caf	intervalized	current active flows	informative
aaf	intervalized	average active flows	informative
paf	intervalized	peak active flows	Check vs max
cfr	intervalized	current flow setup rate	informative
afr	intervalized	average flow setup rate	Check meets expected norms; increasing over time increases load
pfr	intervalized	peak flow setup rate	informative
ctr	intervalized	current traffic rate	informative
atr	intervalized	average traffic rate	Check meets expected norms; increasing over time increases load
ptr	intervalized	peak traffic rate	Check vs max

Table 5 **Tracking ISA Load in the Reporting Interval (Continued)**

Record Name	Type	Description	Load planning use
cpr	intervalized	current packet rate	informative
apr	intervalized	average packet rate	informative
ppr	intervalized	peak packet rate	informative
cds	intervalized	current diverted subscribers	informative
ads	intervalized	average diverted subscribers	informative
pds	intervalized	peak diverted subscribers	Check vs max and expected norms; increasing over time increases load
rfi	intervalized	flows in use	Check vs max and expected norms; increasing over time increases load
rcc	cumulative	ISA capacity cost	Check meets expected norms; increasing over time increases load

The intended deployment model is for this statistic record to be collected by 5620 SAM along with all other AA records and be stored in the 5670 RAM database for subsequent analytics purposes, such as trending charts or setting thresholds of key values. It is recommended that a CRON script be used to export the AA performance record to a storage server for post processing if the 5670 RAM is not deployed:

- If the 5670 RAM is not yet deployed in the network, it is possible to automatically collect the XML accounting files and provide high-level reporting through an XML-to-CSV conversion.
- The simplest approach is to configure a CRON script on the 7x50 SR/ESS to automatically retrieve the CF accounting file (alternatively, any other scripting mechanism with an interval < retention can be used)
- It is recommended that the rollover interval of the file-id policy be modified to 6H or above in order to collect fewer files while keeping the same collection interval.

```
*A:BNG# file type cf2:/script
file copy cf2:/act/*.gz ftp://login:password@IP-ADDRESS/acct/router1/

*A:BNG>config>cron# info
----------------------------------------------
        script "test-ftp-act"
            location "cf2:/script"
            no shutdown
        exit
        action "cron1"
            results "ftp://login:password@IP-ADDRESS/results/router1-result.log"
            script "test-ftp-act"
            no shutdown
```

```
exit
schedule "schedule1"
    interval 36000
    action "cron1"
    no shutdown
exit
```

Note that the interval 36000 is in seconds (10 hours).

With this XML to CSV export mechanism, a spreadsheet can be used by the network engineer to periodically track the ISA resource utilization.

ISA Overload Cut-through

The system can be configured to react to overload based on the weighted-average (WA) queue depth of the shared network port buffer pool from-sub and to-sub. Overload cut-through is typically recommended for use of AA for value-added services where, in the event of overload, the preference is for the ISA to continue to pass packets without AA processing. However, firewall use cases will prefer to drop excess traffic in the event of overload, in which case overload cut-through may not be desired.

In addition to triggering an alarm, further packets sent to the ISA after the WA high-watermark threshold is reached are cut-through immediately by the ISA card without application identification or subscriber policy processing, if the **isa-overload-cut-through** command is enabled.

The WA queue depth is typically configured based on the maximum tolerated delay for the service diverted and the amount of shared buffer space allocated from the IOM (only for IOM3).

AA deployment recommended settings:

- high watermark — 33% of the maximum MBS for all diverted network queues
- low watermark — 5% of the maximum MBS for all diverted network queues

The recommended high and low watermarks assume that the sum of the network port egress queues MBS size is 100% of the shared buffer. If this network queue maximum size is further reduced in the network QoS policy, the watermark values must be adapted proportionally; for example, if the total MBS size cannot exceed 50% of the shared buffer, then the watermark values would be divided by 2 => High Wmark = 16%, Low Wmark = 2%. Adjusting the MBS and the **wa-shared-high-wmark** and **wa-shared-low-wmark** values proportionately ensures that the MBS point (after which discards occur) is above the WA shared high-watermark threshold; otherwise, the ISA will not ever overload if MBS discards are occurring first.

```
A:BNG# configure isa application-assurance-group 1
            isa-overload-cut-through
            qos
                egress
                    from-subscriber
                        wa-shared-high-wmark 16
                        wa-shared-low-wmark 2
                    exit
                    to-subscriber
                        wa-shared-high-wmark 16
                        wa-shared-low-wmark 2
                    exit
                exit
            exit
```

The **show>isa>group** commands can be used to verify that overload cut-though is enabled.

```
*A:BNG>show isa application-assurance-group 1
============================================================
ISA Application-assurance-groups
============================================================
ISA-AA Group Index          : 1
Description                 : (Not Specified)
Subscriber Scale            : residential
WLAN GW Group Index         : N/A
Primary ISA-AA              : 1/2 up/active
Backup ISA-AA               : 2/1 down
Last Active change          : 07/02/2014 12:17:45
Admin State                 : Up
Oper State                  : Up
Diverted FCs                : be
Fail to mode                : fail-to-wire
Partitions                  : enabled
QoS
  Egress from subscriber
    Pool                    : default
      Reserved Cbs          : default
      Slope Policy          : default
    Queue Policy            : default
    Scheduler Policy        :
  Egress to subscriber
    Pool                    : default
      Reserved Cbs          : default
      Slope Policy          : default
```

```
    Queue Policy            : default
    Scheduler Policy        :
Capacity Cost
    High Threshold          : 4294967295
    Low Threshold           : 0
Overload Cut Through        : enabled
Transit Prefix
    Max IPv4 entries        : 0
    Max IPv6 entries        : 0
    Max IPv6 remote entries : 0
HTTP Enrichment
    Max Packet Size         : 1500 octets
===============================================================================
```

To monitor the load status of an ISA, enter the following CLI command.

```
*A:BNG>show application-assurance group 1 status isa 5/1 cpu
========================================
Application-Assurance ISA CPU Utilization
(Test time 993791 uSec)
========================================
Management CPU Usage
----------------------------------------
Name               CPU Time    CPU Usage
                   (uSec)
----------------------------------------
System               14277         1.43%
Management           61101         6.15%
Statistics           69850         7.02%
Idle                848563        85.39%
========================================
Datapath CPU Usage
----------------------------------------
Name              CPU Time     CPU Usage
                  (uSec)
----------------------------------------
System               14277         1.43%
Packet Processing    61101         6.15%
Application ID       69850         7.02%
Idle                848563        85.39%
```

Additionally, the system log files can be used to examine the AA overload history to determine when the overload state was entered and exited. It can be helpful to send AA events to a separate log using the following configuration:

```
    log
        filter 45
            default-action drop
            entry 10
                action forward
                match
                    application eq "application_assurance"
                exit
            exit
        exit
```

```
        log-id 45
            description "application-assurance log"
            filter 45
            from main
            to memory 500
        exit
```

The log files can then be examined to see if overload has occurred, and how frequently. If overload occurs with any regularity, it is a situation that should be addressed. Below is an example of a log file showing AA overload:

```
A:BNG# show log log-id 45
===============================================================================
Event Log 45
===============================================================================
Description : application-assurance log
warning: 13 events dropped from log
Memory Log contents  [size=500   next event=16  (not wrapped)]

15 2014/08/14 17:00:32.66 EST WARNING: APPLICATION_ASSURANCE #4433 Base
"ISA AA Group 1 MDA 5/1 exiting overload cut through processing."

14 2014/08/14 17:00:32.55 EST WARNING: APPLICATION_ASSURANCE #4431 Base
"ISA-AA group 1 MDA 5/1 wa-shared buffer use is less than or equal to 1% in the to-
subscriber direction or corresponding tmnxBsxIsaAaGrpToSbWaSBufOvld notification has
 been disabled."

13 2014/08/14 17:00:32.06 EST WARNING: APPLICATION_ASSURANCE #4432 Base
"ISA AA Group 1 MDA 5/1 entering overload cut through processing."

12 2014/08/14 17:00:32.05 EST WARNING: APPLICATION_ASSURANCE #4430 Base
"ISA-AA group 1 MDA 5/1 wa-
shared buffer use is greater than or equal to 35% in the to-subscriber direction."
```

The primary indicator to look at in CLI statistics for ISA load indication is datapath CPU Usage. Regardless of the configuration and traffic profiles in use, datapath CPU usage gives a consistent indication of whether the ISA is under heavy load (the cause of overload is the inability of the ISA to perform more tasks). The idle datapath time is not proportionate to bandwidth throughput, but if idle datapath CPU usage is under 5%, this indicates an approaching maximum processing load.

At an average datapath use of 95-100% (less than 5% idle) the ISA is creating latency and backpressuring the host IOM egress. It is the best way to know how close to overload the ISA has been. Attempting to examine data throughput statistics such as bit rate, flow setup rate and packet rate to predict overload is not recommended, as these are quite variable under normal circumstances and are not directly correlated to overload. Once in overload, the data statistics (volume, setup rate, etc.) are useful for determining what threshold traps to put in place for the future, but the needed thresholds will always be specific to the live deployment traffic mix and policy configuration.

Below is an example of the status for an ISA that is fully loaded but not yet in overload:

```
*A:BNG>show application-assurance group 1 status isa 5/1 cpu
========================================
Application-Assurance ISA CPU Utilization
========================================

-----------------------------------------------
Management CPU Usage (Test time 999636 uSec)
-----------------------------------------------
Name                  CPU Time     CPU Usage
                      (uSec)
-----------------------------------------------
System                   1540         0.15%
Management                 14        ~0.00%
Statistics             643955        64.42%
ICAP Client               603         0.06%
Idle                   353524        35.37%
-----------------------------------------------

-----------------------------------------------
Datapath CPU Usage   (Test time 999735 uSec)
-----------------------------------------------
Name                  CPU Time     CPU Usage
                      (uSec)
-----------------------------------------------
System                 188374        18.84%
Packet Processing      534203        53.43%
Application ID         277158        27.72%
Idle                        0         0.00%
-----------------------------------------------
```

In this example, 0% idle datapath CPU means the ISA is fully used. When the Datapath CPU Usage Idle average is in the 5-10% range consistently, the ISA should be considered “full”; to add new subscribers, more ISAs are required.

If the excessive traffic condition persists, backpressure from the ISA to the IOM will buffer packets in the egress buffers, and when the egress MBS is exceeded, the ISA host IOM will indicate Diverted discards due to congestion if cut-through is not enabled:

```
*A:BNG>show application-assurance group 1 status detail
===============================================================================
Application-Assurance Status
===============================================================================
Last time change affecting status : 08/12/2014 13:16:15
Number of Active ISAs             : 1
Flows                             : 235754165
Flow Resources In Use             : 12000000
AA Subs Created                   : 14224
AA Subs Deleted                   : 0
AA Subs Modified                  : 1
Seen IP Requests Sent             : 0
Seen IP Requests Dropped          : 0
```

```
-------------------------------------------------------------------------------
                                    Current    Average    Peak
-------------------------------------------------------------------------------
Active Flows                      : 8452434    3786948    10632607
Flow Setup Rate (per second)      : 246578     65104      298677
Traffic Rate (Mbps)               : 33702      13229      35813
Packet Rate (per second)          : 6847697    2466118    6945936
AA-Subs Downloaded                : 14224      13710      14224
Active Subs                       : 14224      9934       14224
-------------------------------------------------------------------------------
                                    Packets                Octets
-------------------------------------------------------------------------------
Diverted traffic                  : 8924242848             5983284952320
Diverted discards                 : 752486                 729147667
    Congestion                    : 752486                 729147667
    Errors                        : 0                      N/A
Entered ISA-AAs                   : 8923417360             5982508976617
Buffered in ISA-AAs               : 57                     19277
Discarded in ISA-AAs              : 0                      0
    Policy                        : 0                      0
    Congestion                    : 0                      0
    Errors                        : 0                      0
Modified in ISA-AAs
    Packet size increased         : 0                      0
    Packet size decreased         : 0                      0
Errors (policy bypass)            : 0                      0
Exited ISA-AAs                    : 8923417303             5982508957340
Returned discards                 : 0                      0
    Congestion                    : 0                      0
    Errors                        : 0                      N/A
Returned traffic                  : 8923285123             5982432640249
===============================================================================
```

ISA Default Subscriber Policy

Default Subscriber Policy — AQP with match criteria not using App-ID or 5-tuple. Match **only** includes traffic direction and/or ASO characteristic and/or subscriber-name.

It is recommended that each ISA be configured with some default subscriber policies that get applied to all subscribers at all times, independent of application flow ID, and even when an ISA is in overload cut-through. These policies protect the ISA resources and provide fairness of resource allocation between subscribers by limiting the ISA resources that can be consumed by a single subscriber. A starting point for the recommended policers is (in all cases, network-specific tuning is recommended):

- Per-subscriber flow rate policer: value more than expected maximum peak per-subscriber rate for active subscribers. The policer protects one subscriber from attacking the network with an excessive flow rate and affecting ISA flow rate resources used by other customers. A typical rate for residential networks could be 100 fps per subscriber.
- Per-subscriber flow count policer: value more than expected maximum per-subscriber flow count for active subscribers. The policer protects one subscriber from consuming excessive flow counts and affecting ISA flow resources used by other customers.
- Downstream bandwidth per subscriber: to a value more than the maximum rate supported by the service, or to less than the maximum per-subscriber capability of the ISA, whichever is lower. For fixed networks, several default policer rates are recommended using a per-sub ASO value for low, medium and large rate ranges set at a rate related to the subscriber access speed. For example, for an FTTH service the per-sub policers could be set at 3 value ranges: below 25Mbps, with another at 100Mbps sub policer for services between 25Mbps and 100Mbps, and another sub-policer for rates between 100Mbps and 1Gbps. The settings for a mobile 3G network rate may be 1Mbps and in an LTE network the rate may be vhcc10Mbs.

For a CLI example of a default subscriber policy, see Application Assurance — App-Profile, ASO and Control Policies.

Conclusion

Any deployment of Application Assurance should include careful capacity planning of the ISA resources, with an appropriate ISA overload strategy, whether for overload cut-through to keep excess traffic flowing, or with a discard policy engineered in the host IOM egress QoS policies.

ISA resource use should be monitored via appropriately configured resource thresholds, events, log files, XML records and show screens to ensure that sufficient ISA resources are available as required.

Application Assurance — HTTP In Browser Notification

This chapter provides information about Application Assurance HTTP in browser notification.

Topics in this chapter include:

Applicability

This chapter is applicable to all 7x50 systems supporting Application Assurance and was tested on SR OS release 13.0.R2.

There are no specific prerequisites for this example.

Overview

Using the 7x50s and Application Assurance, subscribers connected to an operator network can be sent fully customizable on-screen notification messages displayed in a non-disruptive and cost-effective manner through their web browser.

This chapter describes the different options for the operator to customize the notification messages returned to the subscriber using either different HTTP-Notification policies or using the flexible HTTP-URL-PARAM VSA mechanism.

This chapter also describes the additional configuration required with the introduction in SR OS 13.0.R1 of the Notification status monitoring capability allowing the system to notify the subscriber at the next candidate flow instead of waiting for the next notification interval in case the previous notification did not result in a success.

Configuration

The setup comprises of the following elements, see HTTP Notification –Setup:

- 7750 SR + ISA-AA.
- Apache Web Server (delivering notification Javascript and content).
- Subscriber (Desktop/Tablet/Smartphone).
- Authentication, Authorization and Accounting (AAA) for subscriber authentication and Policy Modification.
- Internet Access.

Figure 16 **HTTP Notification –Setup**

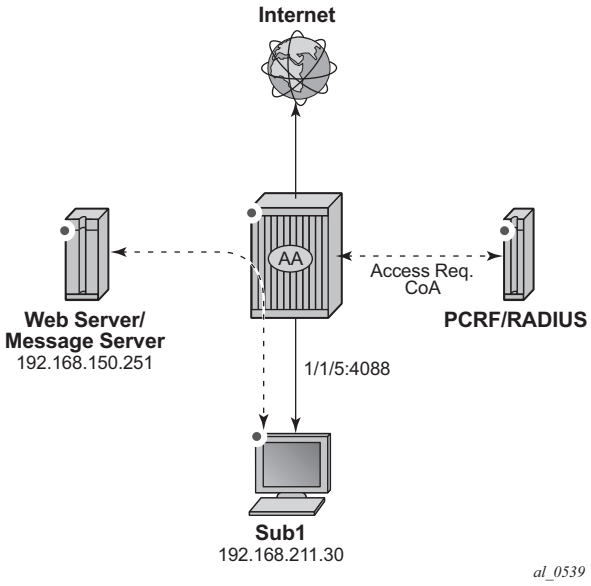

This example describes how to configure HTTP notification to display different notification messages. It demonstrates a simple example in the context of a residential deployment where a message is displayed when the subscriber reaches 80% or 100% of their maximum allowed volume (usage cap).

Figure 17 **Notification Message Example – Quota 80%**

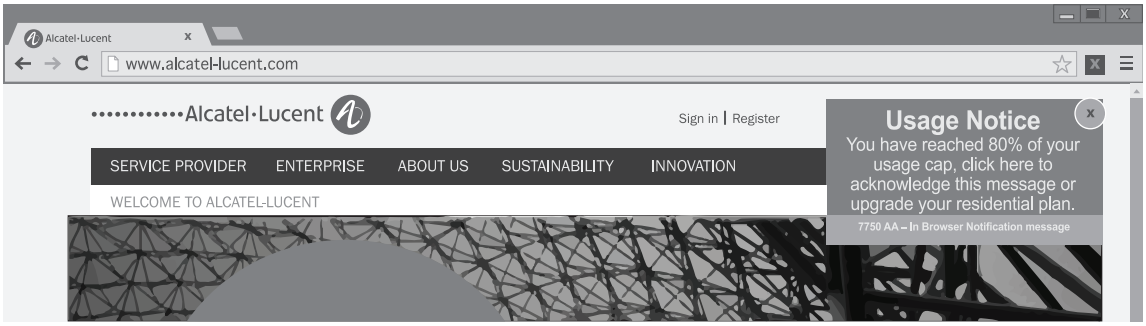

al_0540

In this context the operator has two options:

- Use a dedicated http-notification policy per message type.
- Use a common http-notification policy for any message type together with the newly introduced Http-Url-Param RADIUS VSA.

This example provides configuration examples for both options.

HTTP Notification Policy per Message Type

In this option a dedicated http-notification policy for each notification message is required.

HTTP Notification Policy

Two dedicated HTTP notification policies are used to return a different message to the subscribers when reaching 80% and 100% of their usage cap, the interval in between notifications is set to 15 minutes.

```
configure
    application-assurance group 1
        http-notification "notification-quota-100" create
           description "100% Usage Cap Notification"
           script-url "http://192.168.150.251/In-Browser-Notification/script/quota-
100.js"
           template 1
           interval 15
           no shutdown
        exit

configure
    application-assurance group 1
```

```
        http-notification "notification-quota-80" create
            description "80% Usage Cap Notification"
            script-url "http://192.168.150.251/In-Browser-Notification/script/quota-
80.js"
            template 1
            interval 15
            no shutdown
        exit
```

Notification Status Monitoring

The operator then needs to enable the http-match-all-req feature for any HTTP request sent to the messaging server in order to monitor HTTP notification success and failures. This is done by creating a new application and enabling http-match-all-req within the app-filter.

Success and failure notifications include a specific HTTP encoded URI automatically interpreted as a success or a failure by Application Assurance on a per subscriber basis. If a failure is detected, the system will automatically attempt to notify a new candidate flow instead of waiting for the next notification interval.

```
configure
   application-assurance group 1:1 policy
      application "IBN Messaging Server" create
     app-group "Web"
   exit
   app-filter
     entry 100 create
        expression 1 http-host eq "^192.168.150.251$"
        http-match-all-req
        application "IBN Messaging Server"
        no shutdown
    exit
exit
```

App-Profiles and App-Service-Options

Event based HTTP notifications is enabled by a policy modification triggered via RADIUS or Gx by modifying the subscriber app-profile or using the Application Service Option (ASO) override.

In this implementation of the HTTP notification policy per message type, the following ASO configuration is used:

```
configure
    application-assurance group 1:1 policy
        app-service-option
```

```
            characteristic "quota-message-notification" create
                value "100"
                value "80"
                value "disabled"
                default-value "disabled"
            exit
        exit
        app-profile "1-1/Default" create
            divert
        exit
```

The ASO characteristic **quota-message** values of 100 and 80 enable the App-Qos-Policy (AQP) **notification-quota-100** and **notification-quota-80** as defined below:

```
configure
    application-assurance group 1:1 policy app-qos-policy
        entry 1000 create
            match
                characteristic "quota-message-notification" eq "100"
                application neq "Advertising Statistics"
            exit
            action
                http-notification "notification-quota-100"
            exit
            no shutdown
        exit
        entry 1100 create
            match
                characteristic "quota-message-notification" eq "80"
                application neq "Advertising Statistics"
            exit
            action
                http-notification "notification-quota-80"
            exit
            no shutdown
        exit
```

RADIUS Policy

The following RADIUS CoA message is used to override the ASO characteristic of a residential subscriber so that a notification message can be returned to the subscriber when they reach 80% of their usage cap:

```
NAS-Port-Id = "1/1/5:4088"
Framed-IP-Address = 192.168.211.30
Alc-AA-App-Service-Options = "quota-message-notification=80"
Alc-App-Prof-Str = "1-1/Default"
```

Show Commands

Before the subscriber usage cap limit is reached, and before the RADIUS CoA message is received, the subscriber ASO parameter flag quota-message-notification is set to its default value **disabled** and therefore no App QoS Policy is triggered.

```
A:PE# show application-assurance group 1:1 aa-sub esm "sub1" summary
===============================================================================
Application-Assurance Subscriber Summary (realtime)
===============================================================================
AA-Subscriber           : sub1 (esm)
ISA assigned            : 1/2
App-Profile             : 1-1/Default
App-Profile divert      : Yes
Capacity cost           : 1
Aarp Instance Id        : N/A
HTTP URL Parameters     : (Not Specified)
Last HTTP Notified Time : N/A
-------------------------------------------------------------------------------
Traffic                         Octets              Packets                Flows
-------------------------------------------------------------------------------
... ...
-------------------------------------------------------------------------------
Application Service Options (ASO)
-------------------------------------------------------------------------------
Characteristic                  Value                           Derived from
-------------------------------------------------------------------------------
quota-message-notification      disabled                        default
===============================================================================
```

After the RADIUS CoA message is sent, the subscriber ASO characteristic **quota-message-notification** value is set to **80**, the subscriber-related App QoS Policy entry 1100 now matches for this subscriber:

```
A:PE# show application-assurance group 1:1 aa-sub esm "sub1" summary
===============================================================================
Application-Assurance Subscriber Summary (realtime)
===============================================================================
AA-Subscriber           : sub1 (esm)
ISA assigned            : 1/2
App-Profile             : 1-1/Default
App-Profile divert      : Yes
Capacity cost           : 1
Aarp Instance Id        : N/A
HTTP URL Parameters     : (Not Specified)
Last HTTP Notified Time : N/A
-------------------------------------------------------------------------------
Traffic                         Octets              Packets                Flows
-------------------------------------------------------------------------------
... ...
-------------------------------------------------------------------------------
Application Service Options (ASO)
-------------------------------------------------------------------------------
Characteristic                  Value                           Derived from
-------------------------------------------------------------------------------
```

```
quota-message-notification       80                                  dyn-override
===============================================================================
```

The same command can be used to identify when the last successful subscriber notification occurred, see the Last HTTP Notified Time field:

```
A:PE# show application-assurance group 1:1 aa-sub esm "sub1" summary
===============================================================================
Application-Assurance Subscriber Summary (realtime)
===============================================================================
AA-Subscriber           : sub1 (esm)
ISA assigned            : 1/2
App-Profile             : 1-1/Default
App-Profile divert      : Yes
Capacity cost           : 1
Aarp Instance Id        : N/A
HTTP URL Parameters     : (Not Specified)
Last HTTP Notified Time : 2014/06/24 15:35:49
-------------------------------------------------------------------------------
```

The operator can also identify how many notifications have been sent per http-notification policy per partition:

```
A:PE# show application-assurance group 1 http-notification "notification-quota-80"
===============================================================================
Application Assurance Group 1 HTTP Notification "notification-quota-80"
===============================================================================
Description  : 80% Usage Cap Notification
Template     : 1 - Javascript-url with subId and optional Http-Url-Param
Script URL   : http://192.168.150.251/In-Browser-Notification/script/quota-80.
               js
Admin Status : Up
AQP Ref      : Yes
Interval     : 15 minutes
-------------------------------------------------------------------------------
Group 1:1 Statistics
-------------------------------------------------------------------------------
Notified             : 2                     Succeeded : 2
Criteria Not Matched : 5                     Failed    : 0
===============================================================================
```

The counter Criteria Not Matched is the number of HTTP flows which did not meet the AA ISA flow selection criteria for In Browser Notification. HTTP flow selection is constrained so that only HTTP web pages flows originating from a web browser are targeted, HTTP requests for content such as video or images are not candidate for notification.

HTTP Notification Customization using RADIUS VSA

Instead of using a dedicated HTTP notification policy for every single message type, the operator can return a RADIUS Http-Url-Param VSA at subscriber creation time or via CoA to customize the notification URL using a single policy. This VSA string is automatically appended to the end of the HTTP notification script-url by the 7x50 which can then be used by the web server to decide which notification message to return to the subscriber.

SR OS release supports 1 active HTTP Notification policy per subscriber, 8 different HTTP notification policies per AA ISA group and 1500 different values for the Http-Url-Param VSA. Therefore, using the Http-Url-Param VSA for the customization of the notification is the recommended model to scale the number of notification messages.

For example:

- RADIUS VSA (Alc-AA-Sub-Http-Url-Param): &message=quota80"
- 7750 HTTP Notification configured script-url: http://1.1.1.1/notification.js
- Subscriber HTTP request to the messaging server:

```
http://1.1.1.1/notification.js?SubId=sub1&var=&message=quota80
```

HTTP Notification Policy

A single HTTP notification policy is used to return different notification messages. The interval between notifications is set to 15 minutes.

```
configure
    application-assurance group 1
        http-notification "in-browser-notification" create
            description "Default HTTP Notification Policy"
            script-url "http://192.168.150.251/In-Browser-Notification/script/
                                                    notification-select.php"
            template 1
            interval 15
            no shutdown
        exit
```

Note: This example does not describe the content of the **notification-select.php** file used to parse the URL parameters.

Notification Status Monitoring

The operator then needs to enable the http-match-all-req feature for any HTTP request sent to the messaging server in order to monitor HTTP notification success and failures. This is done by creating a new application and enabling http-match-all-req within the app-filter.

Success and failure notifications include a specific HTTP encoded URI automatically interpreted as a success or a failure by Application Assurance on a per subscriber basis. If a failure is detected, the system will automatically attempt to notify a new candidate flow, instead of waiting for the next notification interval.

```
configure
    application-assurance group 1:1 policy
       application "IBN Messaging Server" create
          app-group "Web"
       exit
    app-filter
       entry 100 create
          expression 1 http-host eq "^192.168.150.251$"
          http-match-all-req
          application "IBN Messaging Server"
          no shutdown
       exit
    exit
```

App-Profile and App-Service-Options

Similar to the previous example, HTTP notifications are enabled per subscriber using RADIUS or Gx by modifying the subscriber app-profile or using ASO override.

The following ASO configuration is used:

```
configure
    application-assurance group 1:1 policy
        app-service-option
            characteristic "in-browser-notification"
                value "enabled"
                value "disabled"
                default-value "disabled"
        exit
```

The ASO characteristic in-browser-notification value **enabled** is used to enable the app-qos-policy matching the http-notification policy in-browser-notification as shown below:

```
configure
```

```
application-assurance group 1:1 policy app-qos-policy
    entry 1300 create
        match
            characteristic "in-browser-notification" eq "enabled"
            application neq "Advertising Statistics"
        exit
        action
            http-notification "in-browser-notification"
        exit
        no shutdown
```

RADIUS Policy

The following RADIUS CoA message is used to modify the ASO characteristic of a residential subscriber and assign a specific Http-Url-Param VSA. The **in-browser-notification** ASO characteristic with value **enabled** is dynamically assigned to the subscriber along with the **Http-Url-Param &message=quota80**:

```
NAS-Port-Id = "1/1/5:4088"
Framed-IP-Address = 192.168.211.30
Alc-AA-App-Service-Options = "in-browser-notification=enabled"
Alc-AA-Sub-Http-Url-Param = "&message=quota80"
Alc-App-Prof-Str = "1-1/Default"
```

The subscriber HTTP request to the messaging server has the following format and includes the Http-Url-Param value as an argument of the URL:

```
http://192.168.150.251/In-Browser-Notification/script/notification-
select.php?SubId=sub1&var=&message=quota80
```

The web server can now use the parameter value to make a decision to return a suitable notification message related to the subscriber usage cap.

Show Commands

Both the **in-browser-notification** ASO characteristic with value **enabled** and the HTTP-Url-Param VSA can be shown as follows:

```
A:PE# show application-assurance group 1:1 aa-
sub esm "sub1" summary
===============================================================================
Application-Assurance Subscriber Summary (realtime)
===============================================================================
AA-Subscriber           : sub1 (esm)
ISA assigned            : 1/2
App-Profile             : 1-1/Default
App-Profile divert      : Yes
```

```
Capacity cost           : 1
Aarp Instance Id        : N/A
HTTP URL Parameters     : &message=quota80
Last HTTP Notified Time : N/A
-------------------------------------------------------------------------------
Traffic                        Octets              Packets                Flows
-------------------------------------------------------------------------------
... ...
-------------------------------------------------------------------------------
Application Service Options (ASO)
-------------------------------------------------------------------------------
Characteristic                  Value                            Derived from
-------------------------------------------------------------------------------
in-browser-notification         enabled                          dyn-override
quota-message-notification      disabled                         default
===============================================================================
```

The operator can also display the HTTP URL parameters VSA currently in use, per AA ISA group:

```
A:PE## tools dump application-assurance group 1 http-url-param-list
-------------------------------------------------------------------------------
Application-Assurance Subscriber HTTP URL parameters for Group 1:
-------------------------------------------------------------------------------
============================================
Http Url Parameter                 Sub Usage
--------------------------------------------
"&message=quota80"                 1
============================================
Total entries displayed 1
```

Conclusion

This chapter, intended for Application Assurance (AA) network architects and engineers, provides two implementation options for configuring and deploying HTTP In Browser Notification. It also explains how to take advantage of the Http-Url-Param RADIUS VSA to flexibly define various messaging campaigns using a common AA notification policy.

Application Assurance — Local URL List Filtering

This chapter provides information about the Application Assurance local URL list filtering.

Topics in this chapter include:

Applicability

This chapter is applicable to all 7x50 systems supporting Application Assurance and was tested on SR OS release 13.0.R2.

There are no specific pre-requisites for this feature.

Overview

The local URL-list filtering capability provided by Application Assurance prevents subscribers from accessing URLs that match a list of restricted URLs stored locally in the 7x50 system compact flash (CF).

This capability assists service providers to comply with regulatory requirements for network-wide URL filtering policies, such as:

- Court-ordered URL takedown
- Child protection
- Government-mandated URL takedown list

The 7x50 uses the Application Assurance capabilities to extract the URL from a subscriber HTTP/HTTPS request and compare it to the list of URLs contained in the local file. If a match occurs, the subscriber request is redirected to a preconfigured web server landing page, typically describing why the access to this resource was denied.

The system supports both unencrypted and OpenSSL Triple Data Encryption Standard (3DES) encrypted file formats to protect the content of the list.

URL-List Update

The system supports a flexible mechanism to upgrade the URL list automatically, using either Cron or the 5620 SAM, to comply with the regulatory requirements for list upgrade frequency.

HTTP/HTTPS Filtering

Each HTTP request within a TCP flow is analyzed and filtered. For HTTPS traffic, the system extracts the domain name information contained in the Transport Layer Security (TLS) Server Name Indication (SNI).

Setup Details

The setup consists of the following elements, as shown in Figure 18:

- 7x50 + ISA-AA
- Web Server (redirect landing page)
- FTP Server (source for the URL-list file)
- Subscriber (desktop/laptop/tablet/smart phone)
- Internet access
- Optional: AAA for subscriber authentication and policy modification

Figure 18 **Local URL-List Filtering Setup**

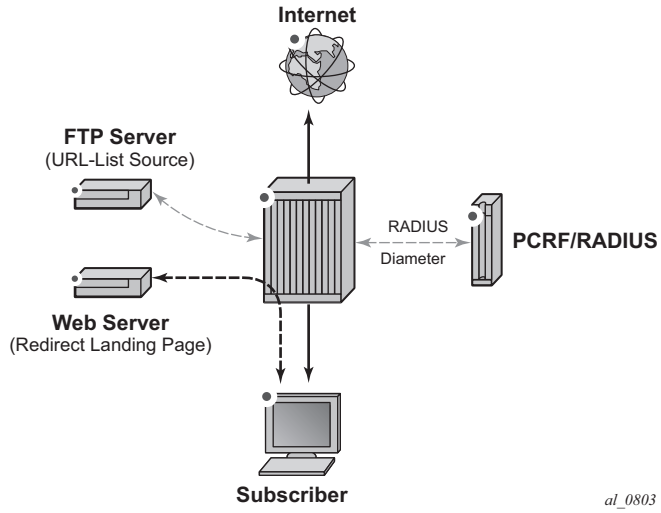

This chapter is written in the context of a residential or WiFi deployment. However, local URL-list filtering is also applicable to business VPN services.

Configuration

To configure the system for local URL-list filtering, the operator needs to:

- Create a URL-list policy referencing a valid URL-list file located on the system compact flash
- Create a URL-filter policy for local filtering by referencing the URL-list policy previously created
- Create an App-QoS-Policy (AQP) to apply this url-filter policy

URL-List Policy and URL-Filter Policy

In the following example, two dedicated URL-list and URL-filter policies show URL filtering based on a plain text file and an encrypted file:

```
configure
    application-assurance group 1
        url-list "blacklist1-encrypted" create
            description "Demo URL Filtering List - Encrypted File"
```

```
            decrypt-key "ON3HU2GFPHmpOHwWbSGw/zdM4iuxzySpqS7pw/
u3qIcuG4mABmrhc." hash2
            file "cf3:\aa-url-list\url-list1.encrypted"
            no shutdown
        exit
        url-filter "local-filter-list1-encrypted" create
            default-action allow
            http-redirect "redirect-blacklist"
            local-filtering
                url-list "blacklist1-encrypted"
            exit
            no shutdown
        exit

configure
    application-assurance group 1
        url-list "blacklist1-plaintext" create
            description "Demo URL Filtering List - Plaintext File"
            file "cf3:\aa-url-list\url-list1-plaintext.txt"
            no shutdown
        exit
        url-filter "local-filter-list1-txt" create
            default-action allow
            http-redirect "redirect-blacklist"
            local-filtering
                url-list "blacklist1-plaintext"
            exit
            no shutdown
        exit
```

In the preceding example, both URL-filter policies are defined using default-action allow. The default action is used in case the file could not be loaded by the system, either at boot time or the first time the URL-list file was configured in the system. Possible causes are, for example:

- File corrupted, compact flash corrupted
- Incorrect file encryption format or password
- Wrong URL format in the file
- Too many URLs in the file

Operators should always use default-action allow when configuring the URL-filter policy associated with a URL-list file because the file or the CF may be corrupted, in which case the system logs an error and a trap is raised.

Note that if a valid URL-list file was previously in use, and an invalid file is uploaded and the URL-list policy upgraded using this file, then the system will continue using the previous list.

HTTP-Redirect Policy

Both URL-filter policies defined in the preceding example refer to the following http-redirect policy; subscribers accessing a URL from the URL-list file are redirected to the following landing page:

```
configure
    application-assurance group 1
        http-redirect "redirect-blacklist" create
            description "Redirect for Local List URL Filtering"
            template 5
            tcp-client-reset
            redirect-url "http://172.16.70.100/Redirect/redirect-
blacklist.html?Request
                         edURL=$URL"
            no shutdown
        exit
```

URL-List File

File Format

The system supports the following format for the URLs contained in the URL-list file:

- URLs without the HTTP keyword. For example:

  ```
  www.domain.com/path
  ```

- URLs with the HTTP keyword. For example:

  ```
  http://www.domain.com/path
  ```

- Comment lines starting with the number sign character (#). For example:

  ```
  # This is a comment line
  ```

- Printable ASCII characters. URLs using non-printable ASCII characters are percent-encoded by the web browser automatically and, therefore, need to be percent-encoded in the URL-list file.

File Encryption

OpenSSL triple DES -nosalt is the supported encryption format. Files can be encrypted offline on a server using the following command:

```
openssl des3 -nosalt -in <input.txt> -out <output.enc>
```

List Upgrade

The URL-list file can be upgraded using the **admin** command:

```
A:BNG# admin application-assurance group 1 url-list "blacklist1-plaintext" upgrade
```

The upgrade result is logged in the system log-id 99:

```
A:BNG# show log log-id 99
===============================================================================
Event Log 99
===============================================================================
Description : Default System Log
Memory Log contents  [size=500   next event=72  (not wrapped)]
71 2015/07/07 13:09:25.01 EST MINOR: APPLICATION_ASSURANCE #4446 Base url-
list success
"URL list "blacklist1-plaintext" in ISA-
AA group 1 has been updated. There are 3 entries in the URL list."
```

App-Profiles and App-Service-Options

Application Assurance policies can be selectively applied to specific AA subscribers by modifying the app-profile assigned to the subscriber or using Application Service Option (ASO) override. See SR OS User Guide for more information about modifying the app-profile or ASO assigned to AA subscribers (RADIUS, Gx, Override).

In this example, the following ASO configuration is used:

```
configure
    application-assurance group 1:1 policy
        app-service-options
            characteristic "local-list-filtering" create
                value "no"
                value "yes-encrypted"
                value "yes-plaintext"
                default-value "no"
            exit
        exit
        app-profile "1-1/Default" create
            divert
        exit
```

The ASO characteristic local-list-filtering value of yes-encrypted and yes-plaintext enable the AQP entry 210 and 220 in the example:

```
configure
```

```
application-assurance group 1:1 policy app-qos-policy
    entry 210 create
        match
            characteristic "local-list-filtering" eq "yes-encrypted"
        exit
        action
            url-filter "local-filter-list1-encrypted"
        exit
        no shutdown
    exit
    entry 220 create
        match
            characteristic "local-list-filtering" eq "yes-plaintext"
        exit
        action
            url-filter "local-filter-list1-txt"
        exit
        no shutdown
    exit
```

If the url-filter policy needs to be applied to 100% of the subscribers in the network, it is also possible to remove the ASO match criteria.

Show Commands

url-list

The status of the URL list can be shown in the CLI. The url-list show command provides basic admin and operational status, as well as the number of URLs in the list. The command also provides reasons for any possible issue related to loading the list, as well as the last successfully deployed file and the last upgrade attempt. Therefore, the operator can determine whether the latest version of the file is currently in use or if an error occurred when trying to upgrade the list.

Show command output:

```
Label                    Description
Admin                    Status [Up | Down] - Administrative status of the url-list
Oper Status              [Up | Down] - Operational status of the url-list
Oper Flags               [admin-down | file-does-not-exist |invalid-file-
format | too
                          -many-urls| switch-over-error]
File Deployed to ISA      [Yes | No] -
 This flag describes if the file located in the
                          compact flash is the one deployed in the ISA, in the event
                           the file is overwritten and before the admin upgrade comm
and
                          is used this flag will display "No".
Upgrade Statistics
```

```
Last Success              Last time the list was successfully upgraded
File Name                 File name for the last successful upgrade
URL Entries               Number of URLs loaded at the last success
Blank/Comment Lines       Number of blank or commented out lines
Last Attempt              Last time the operator tried to upgrade the list
Result                    Success | Failure. Result of the last upgrade
File Name                 File name for the last upgrade attempt

A:BNG# show application-assurance group 1 url-list "blacklist1-plaintext"
===============================================================================
Application Assurance Group 1 url-list "blacklist1-plaintext"
===============================================================================
Description              : Demo URL Filtering List - Plaintext File
Admin Status             : Up
Oper Status              : Up
Oper Flags               : <none>
File deployed to ISAs    : Yes
-------------------------------------------------------------------------------
Upgrade Statistics
-------------------------------------------------------------------------------
Last Success             : 07/07/2015 11:56:01
  Deployed
    File Name            : cf3:\aa-url-list\url-list1-plaintext.txt
    URL Entries          : 3
    Blank/Comment Lines  : 2
Last Attempt             : 07/07/2015 11:56:01
  Result                 : Success
    File Name            : cf3:\aa-url-list\url-list1-plaintext.txt
===============================================================================
```

url-filter

The **url-filter** show command provides its operational and admin status, as well as actions taken, such as the number of redirects. With URL list filtering, using a **default-action** set to **allow**, the only counters increasing are **allow**, **redirect**, and **default**.

```
A:BNG# show application-assurance group 1 url-filter "local-filter-list1-txt"
===============================================================================
Application Assurance Group 1 URL Filter "local-filter-list1-txt"
===============================================================================
Description            : (Not Specified)
Admin Status           : Up
Oper Status            : Up
Oper Flags             : <none>

Default Action         : allow
HTTP Request Filtering : all
HTTP Redirect          : redirect-blacklist
AQP Referenced         : Yes

Local Filter URL List  : blacklist1-plaintext
    Admin Status       : Up
    Oper Status        : Up
```

```
    Number of URLs     : 3
    Last Update Time   : 07/07/2015 11:56:01

-------------------------------------------------------------------------------
Total HTTP Filtering Stats
-------------------------------------------------------------------------------
HTTP Requests   : 26
HTTP Req Errors : 0

HTTP Response Actions
    Allow       : 23
    Block       : 0
    Redirect    : 3
    Default     : 0
===============================================================================
```

http-redirect

The **http-redirect** show command provides more information about how the traffic was blocked; for example, it differentiates TCP client reset used for HTTPS from regular redirect used for HTTP traffic.

```
A:BNG# show application-assurance group 1 http-redirect "redirect-blacklist"
===============================================================================
Application Assurance Group 1 HTTP Redirect redirect-blacklist
===============================================================================
Description        : Redirect for Local List URL Filtering
Template           : 5
                   : Redirect supporting macro substitution using HTTP 302
Redirect URL       : http://172.16.70.100/Redirect/redirect-blacklist.
                     html?RequestedURL=$URL
Admin Status       : Up
AQP Ref            : No
-------------------------------------------------------------------------------
Summary Statistics
-------------------------------------------------------------------------------
Grp:Part             Redirects           Client Resets               Redirects
                          Sent                    Sent                Not Sent
-------------------------------------------------------------------------------
1:1                          2                       1                       0
-------------------------------------------------------------------------------
Total                        2                       1                       0
-------------------------------------------------------------------------------
===============================================================================
```

Conclusion

This chapter, intended for Application Assurance (AA) network architects and engineers, provides two examples for deploying URL-list filtering as well as upgrading the list and displaying its statistics.

Application Assurance — Security Gateway Stateful Firewall

This chapter provides information about Application Assurance Security gateway stateful firewall.

Topics in this chapter include:

Applicability

This chapter is applicable to all 7750 SR/SR-c and 7450 ESS chassis supporting Application Assurance (AA).

The configuration was tested on release 13.0.R2.

Overview

The SR OS 13.0.R1 AA stateful firewall feature runs on AA-ISA and extends application-level analysis to provide an in-line stateful service, integrated within the Security Gateway (SeGW). The feature provides protection for mobile infrastructure; Mobility Management Entities (MMEs), Serving Gateways (SGWs), and Network Management Systems (NMSs), against attacks from compromised base stations, evolved NodeBs (eNBs), or Femto Access Points (FAPs). AA stateful packet filtering, combined with AA L7 classification and control, provides advanced, next-generation firewall functionality. Using stateful packet filtering, the AA FW not only inspects packets at layers 3 to 7, but also monitors the connection state.

AA FW deployed within a SeGW in ultra-broadband access networks (3G/4G/Femto) provides back-end core network security protection, as per Figure 1. AA FW offers protection for the following 3rd Generation Partnership Project (3GPP) defined interfaces:

1. S1-MME

2. S1-U
3. Operations, Administration and Maintenance (OAM)

Figure 19 **LTE SeGW Firewall Deployment**

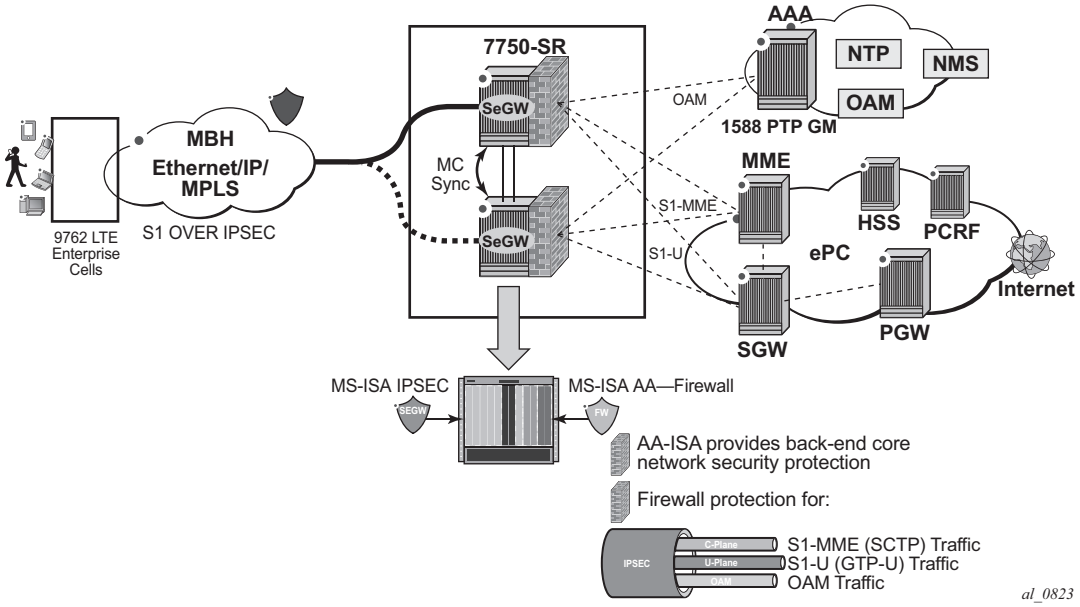

Similarly, the SeGW architecture for Femto deployment is based on two 7750 SR systems terminating the mobile backhaul side (front-end and connecting to, for example, a 9365 Base Station Router Gateway (BSR-GW) and other network elements of the packet core (back-end), as per Figure 20:

Figure 20 **SeGW in Small Cells Architecture**

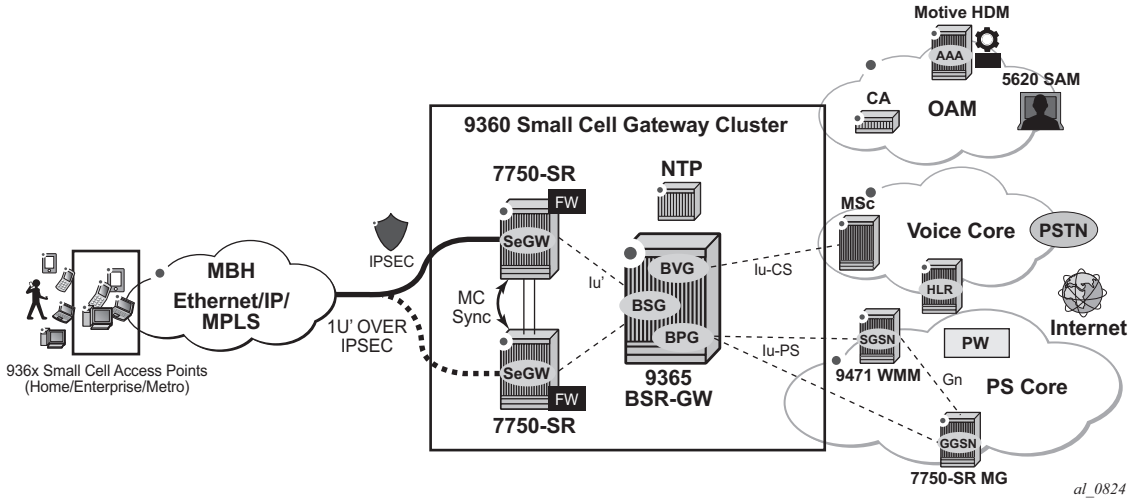

The two SeGWs run in stateful redundant mode: upon partial or total failure of the active SeGW for a set of IPSec tunnels, the other SeGW takes over without terminating the IPSec tunnels, providing hitless failover.

In addition to MS-ISA hardware dedicated to the IPSec function, each SeGW supports one or more additional MS-ISAs running AA to provide firewall capabilities. The firewall rules protect the BSR as well as the BSR-GW and packet core network elements (NEs) from malicious attacks or unauthorized traffic.

The objective of this chapter is to describe the required configuration within AA-ISA in order to enable AA FW and protection for S1-MME, S1-U, and OAM traffic. Basic knowledge of AA-ISA diversion configuration is assumed.

S1-MME Traffic Protection

The purpose of AA FW in this deployment is to protect the MME infrastructure against an attack from a compromised eNB or FAP. Network flooding attacks, malformed packets, and port scans are examples of denial of service (DoS) attacks that can be carried out using a compromised eNB or FAP.

AA FW provides inspection of the Stream Control Transmission Protocol (SCTP) used to communicate to the MME. Such inspection includes checking for SCTP payload protocol IDs (PPIDs), source /destination ports, SCTP chunk validation, and malformed SCTP packets (such as checksum validation). In addition, the operator can configure DoS flooding rules, such as policers to limit the bandwidth and/or flow counts of SCTP traffic.

S1-U Traffic Protection

The purpose of AA FW in this deployment is to protect the SGW infrastructure against an attack from a compromised eNB or FAP. AA FW supports protection against:

- malformed GPRS Tunneling Protocol User plane (GTP-U) packet attacks

  Checking packet sanity, which include GTP-U mandatory, optional, and extension header checks, as well as checks for invalid reserved information elements (IE) and missing IEs.
- unsupported GTP messages

  Filtering messages based on message type and/or message length.
- flooding attacks

Shaping GTP traffic bandwidth, which limits the GTP-U bandwidth that a FAP can send to the core (SGW).

Limiting GTP tunnels, which limits the number of concurrent GTP tunnels and/or setup rate of these tunnels from a FAP to the core network.

To prevent the shared resources of bandwidth and the SGW processor from being consumed by an attacker, Nokia recommends the GTP flow rate limiting configuration.

- IP fragmentation-based attacks

  Applying various drop rules for IP fragmentation of GTP messages.

OAM Traffic Protection

The purpose of AA FW protection in this deployment is to protect against any abuse of OAM network resources, such as NMS.

Network flooding attacks, malformed packets, and port scans are examples of such attacks that can be carried out using a compromised eNB or FAP.

See the configuration described in the Application Assurance — Stateful Firewall chapter for this context of OAM protection in SeGW.

Configuration

AA-ISA Application QoS Policies (AQPs) are enhanced in Release 13.0.R1 with several new AQP actions that provide SCTP and GTP filtering functionality. As with all AQPs, these actions have partition-level scope, which allows different FW policies to be implemented by using AA partition concepts within the same AA-ISA.

The configuration topology in Figure 21 shows the SeGW FW functionality of S1-U and S1-MME interfaces. Geo-redundancy, which is a very common deployment option, is not described in this here because it is described in the Multi-Chassis IPSec Redundancy chapter.

Figure 21 **Configuration Topology**

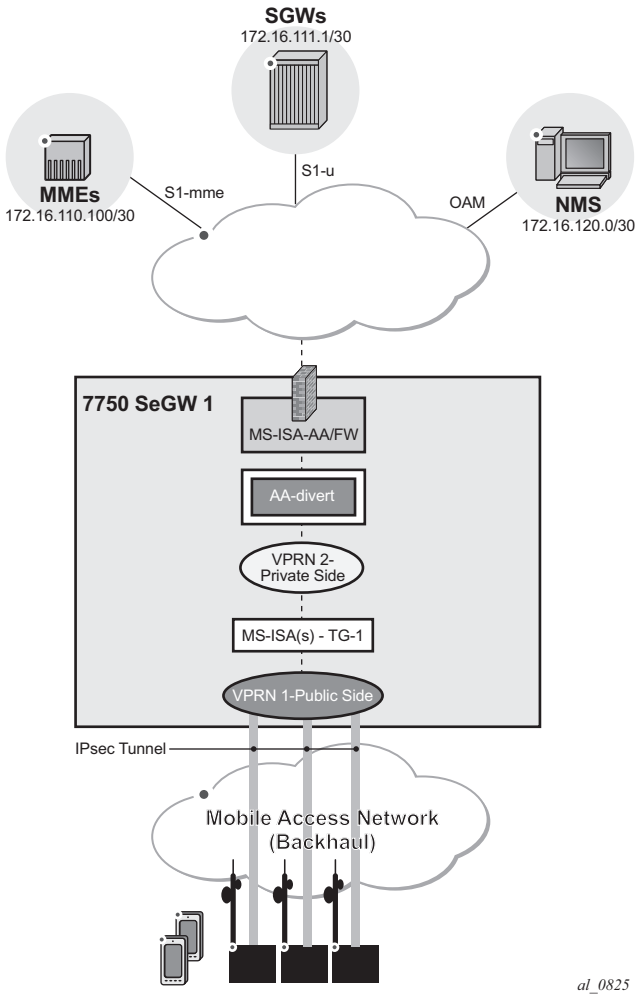

Pre-Setup Requirements

Configure tunnel ISAs with optional multi-chassis redundancy. See the Multi-Chassis IPSec Redundancy chapter for more information.

Step 1. Divert AA traffic and apply basic firewall rules.

Step 1.1. Divert private VPRN traffic into AA-ISA with AA multi-chassis redundancy.

This step is required for any of the configurations in Steps 2, 3 or 4.

There is no dependency between Steps 2, 3 or 4.

In this example, one private VPRN is used for all traffic to/from eNBs. In some small cell deployments, eNB traffic is split into three different VPRNs: one for control (S1-MME), one for management (OAM), and one for bearer traffic (S1-U GTP-U). In that case, each of these VPRNs needs to be diverted into AA-ISA in order to provide firewall protection.

First, define an application profile and transit IP policy, such as:

```
*A:7750-1>config>app-assure>group$ info
----------------------------------------------
            policy
                begin
                app-profile "default" create
                    description " App profile that applies to the whole SAP"
                    divert
                exit
                commit
            exit
            transit-ip-policy 1 create
                description "Per eNB-IP Sub policy"
                detect-seen-ip
                transit-auto-create
                    no shutdown
                exit
```

Then, apply these policies to the SAP on the private side of the IPSec tunnel ISA:

```
*A:7750-1>config>service>vprn>if# info
----------------------------------------------
                sap tunnel-1.private:1 create
                    transit-policy ip 1
                    app-profile "default"
                exit
```

This configuration achieves:

- Traffic to/from the IPSec tunnel ISA private SAP is diverted to AA-ISA for the purpose of FW protection,
- Within AA-ISA, the diverted SAP is treated as a parent SAP. That is, instead of treating the whole SAP as a single subscriber, subscribers are auto-created within this SAP based on the IP address of the eNBs.

Step 1.2. Protect against malformed packets.

In firewall deployments, it is recommended that overload-drop, error-drop, and fragment-drop (all) are enabled within the default sub-policy, as shown in the example below:

- **overload-drop** ensures that AA-ISA, when overloaded, drops the excess traffic instead of allowing it through, without applying firewall rules.
- **error-drop** ensures that AA-ISA drops malformed IP packets.

– **fragment-drop (all)** because many network DoS attacks use IP fragmentation to initiate attacks, the operator has the option to drop all fragmented traffic, drop out-of-order fragments only, or allow fragments through. Allowing fragments through is not recommended for firewall deployments.

```
*A:7750-1>config>app-assure>group>policy# app-qos-policy
*A:7750-1>config>app-assure>group>policy>aqp# info
----------------------------------------------
                    entry 500 create
                        description "apply SeGW session filter rules"
                        match
                            traffic-direction subscriber-to-network
                        exit
                        action
                            overload-drop
                            error-drop
                            fragment-drop all
                        exit
                        no shutdown
                    exit
                exit
----------------------------------------------
*A:7750-1>config>app-assure>group>policy#
```

Step 1.3. Limit total traffic from any eNB.

It is recommended that a total limit be placed on how much bandwidth and how many flows an eNB or FAP can generate toward the network, regardless of the type of traffic.

The limit values are a function of the number of end devices that are served by the eNB or FAP, plus some additional margin:

```
*A:7750-1>config>app-assure>group# info
----------------------------------------------
            policer "limit_eNBs_total_Flows" type flow-count-
limit granularity subscriber
                create
                flow-count 1000
            exit
            policer "limit_eNBs_total_bw" type single-bucket-
bandwidth granularity sub
                scriber create
                rate 500
                mbs 500
            exit
----------------------------------------------
*A:7750-1>config>app-assure>group#
```

Note: If the traffic from eNB or FAP is separated into different private SAPs, based on traffic type (S1-AP, S1-U, or OAM), as with some deployment topologies, then the policing limit value is dependent on the SAP traffic type as well as the number of end devices. See policing limit settings in Steps 2 and 3.

Apply the configured policers as actions from within the default sub-policy AQP entry:

```
*A:7750-1>config>app-assure>group>policy# app-qos-policy entry 500
*A:7750-1>config>app-assure>group>policy>aqp>entry>action# flow-count-
limit "limit_eNBs_total_Flows"
*A:7750-1>config>app-assure>group>policy>aqp>entry>action# bandwidth-
policer "limit_eNBs_total_bw"
*A:7750-1>config>app-assure>group>policy>aqp>entry# info
----------------------------------------------
                        description "apply SeGW session filter rules"
                        match
                            traffic-direction subscriber-to-network
                        exit
                        action
                            bandwidth-policer "limit_eNBs_total_bw"
                            flow-count-limit "limit_eNBs_total_Flows"
                            session-filter "SeGW_FW"
                            overload-drop
                            error-drop
                            fragment-drop all
                        exit
                        no shutdown
----------------------------------------------
*A:7750-1>config>app-assure>group>policy>aqp>entry#
```

Note: All of the above listed actions use the traffic direction of subscriber-to-network. That is, they are not applied to traffic in the other direction (downstream) because the purpose of the firewall is to protect the network resources from upstream traffic coming from compromised eNBs or FAPs.

Step 2. Configure AA-ISA to provide firewall protection to protect MMEs (S1-AP traffic).

Step 2.1. Create IP AA lists.

First, create an AA IP prefix list that contains eNB IP addresses or range of addresses:

```
*A:7750-1>config>app-assure# group 1:1
*A:7750-1>config>app-assure>group# ip-prefix-list "ALL_eNBs" create
*A:7750-1>config>app-assure>group>pfx>$ description "eNodeB subnet"
*A:7750-1>config>app-assure>group>pfx>$ prefix 172.16.100.0/24
*A:7750-1>config>app-assure>group>pfx>$ exit
```

Next, optionally create an AA IP list that contains MME IP addresses (in case there are more than one):

```
*A:7750-1>config>app-assure>group# ip-prefix-list "MMEs" create
*A:7750-1>config>app-assure>group>pfx>$ description "MME(s) subnet"
*A:7750-1>config>app-assure>group>pfx>$ prefix 172.16.110.100/30
*A:7750-1>config>app-assure>group>pfx>$ exit
```

After the above lists are created, they can be referenced and used in AA FW rules using session filters and AQPs.

Step 2.2. Allow only SCTP traffic towards MMEs — No port scanning.

A basic setup creates session-filter rules that will only allow SCTP traffic between eNBs and MMEs.

```
*A:7750-1>config>app-assure>group# session-filter "SeGW_FW" create
*A:7750-1>config>app-assure>group>sess-fltr$ default-action deny
*A:7750-1>config>app-assure>group>sess-fltr$ entry 1 create
*A:7750-1>config>app-assure>group>sess-fltr>entry$ description "allow SCTP to MM Es"
*A:7750-1>config>app-assure>group>sess-fltr>entry$ match
*A:7750-1>config>app-assure>group>sess-fltr>entry>match$ ip-protocol-num "sctp"
*A:7750-1>config>app-assure>group>sess-fltr>entry>match$ src-ip ip-prefix-
list "ALL_eNBs"
*A:7750-1>config>app-assure>group>sess-fltr>entry>match$ dst-ip ip-prefix-
list "MMEs"
*A:7750-1>config>app-assure>group>sess-fltr>entry>match$ dst-port eq 6005
*A:7750-1>config>app-assure>group>sess-fltr>entry>match$ exit
*A:7750-1>config>app-assure>group>sess-fltr>entry$ action permit
*A:7750-1>config>app-assure>group>sess-fltr>entry$ exit
```

Note: In the above configuration, SCTP traffic on MMEs is assumed to be running on port 6005.

Next, the newly created session filter needs to be referenced from a default sub-policy AQP action, as follows:

```
*A:7750-1>config>app-assure>group>policy>aqp# info
----------------------------------------------
                    entry 500 create
                        description "apply SeGW session filter rules"
                        match
                            traffic-direction subscriber-to-network
                        exit
                        action
                            session-filter "SeGW_FW"
                            overload-drop
                            error-drop
                            fragment-drop all
                        exit
                        no shutdown
                    exit
                exit
----------------------------------------------
*A:7750-1>config>app-assure>group>policy#
```

Using traffic direction **subscriber-to-network** in the above AQP entry achieves two objectives:

1. Protects MMEs by allowing only SCTP traffic to be initiated from eNB subnets toward MMEs. Port scanning toward MME is blocked.

2. Allows MMEs to have full access to eNBs.

Note: It is important that an AQP, containing a session-filter action, does not contain any matching condition other than ASOs, traffic direction, or subscriber ID. Subscriber ID is not applicable in this deployment use-case.

Step 2.3. DoS protection: Limit the number of SCTP flows from eNBs.

In this step, the operator configures a flow count policer to limit the number of SCTP flows that an eNB can generate toward the MMEs. This protects the MMEs against a compromised eNB trying to set up many SCTP flows.

```
*A:7750-1# configure application-assurance group 1
*A:7750-1>config>app-assure>group# policer "sctp_flow_count" type flow-count-
limit granularity subscriber create
*A:7750-1>config>app-assure>group>policer$ flow-count 2
*A:7750-1>config>app-assure>group>policer$ exit
```

In the above configuration, an eNB or FAP can have up to two flows at a time. In practice, there should only be one SCTP session, one flow in each direction, per eNB-MME pair. The above example uses two flows to leave a margin in case a second, backup, MME needs to communicate with the eNB, while still providing enough protection.

Add the defined policer as a **flow-count-limit** as an AQP action, as follows:

```
A:7750-1>config>app-assure>group>policy>aqp# entry 100
A:7750-1>config>app-assure>group>policy>aqp>entry$ info
----------------------------------------------
                        description "limit SCTP traffic"
                        match
                            traffic-direction subscriber-to-network
                            ip-protocol-num eq sctp
                        exit
                        action
                            flow-count-limit "sctp_flow_count"
                        exit
                        no shutdown
----------------------------------------------
A:7750-1>config>app-assure>group>policy>aqp>entry$
```

Step 2.3.1. Configure an AA FW events log.

It is sometimes advisable to configure a log that captures events related to various AA FW actions. Due to the limited size of the log and the large amount of traffic AA can handle, consider the usefulness of the information in the log when:

- debugging a configuration
- testing a configuration in a staged environment
- capturing infrequent actions

To configure a log:

```
*A:7750-1# configure application-assurance group 1:1
*A:7750-1>config>app-assure>group# event-log "FW_drops_log" create
*A:7750-1>config>app-assure>group>evt-log$ buffer-type circular
*A:7750-1>config>app-assure>group>evt-log$ max-entries 100000
*A:7750-1>config>app-assure>group>evt-log$ no shutdown
*A:7750-1>config>app-assure>group>evt-log$ exit
*A:7750-1>config>app-assure>group# info
----------------------------------------------
                ---snipped---
            event-log "FW_drops_log" create
                buffer-type circular
                max-entries 100000
                no shutdown
            exit
```

To reference the configured log from within the deny action of the session filter:

```
*A:7750-1>config>app-assure>group>sess-fltr# info
----------------------------------------------
                default-action deny event-log "FW_drops_log"
                entry 1 create
                    description "allow SCTP to MMEs"
                    match
                        ip-protocol-num sctp
                        src-ip ip-prefix-list "ALL_eNBs"
                        dst-ip ip-prefix-list "MMEs"
                    exit
                    action permit
                exit
----------------------------------------------
*A:7750-1>config>app-assure>group>sess-fltr#
```

To view the log:

```
*A:7750-1# tools dump application-assurance group 1:1 event-
log "FW_drops_log" isa 1/2
===============================================
Application-Assurance event-log "FW_drops_log"
Current Time:          "06/10/2015 22:45:30" (UTC)
  group[:partition]:   1:1
  isa:                 1/2
  admin state:         no shutdown
  buffer-type:         circular
  max-entries:         100000
===============================================
Event-
source
                     Action        SubType     SubName                              Di
rection Src-ip                                    Dst-
ip                                 Ip-protocol Src-port Dst-port Timestamp

Total Records:   0
===============================================
*A:7750-1#
```

To clear all the entries within the specified log:

```
*A:7750-1# clear application-assurance group 1:1 event-log "FW_drops_log"
```

Step 2.4. DoS protection: Limit the SCTP bandwidth from eNB

Similar to the previous step, the operator configures a flow bandwidth policer to limit the amount of SCTP traffic that an eNB can generate toward the MMEs. This protects the MMEs against a compromised eNB trying to flood the MMEs.

```
*A:7750-1# configure application-assurance group 1
*A:7750-1>config>app-assure>group# info
----------------------------------------------
---snipped---
            policer "sctp_bw_limit" type single-bucket-
bandwidth granularity subscriber create
                rate 30
                mbs 10
            exit
---snipped---
            exit
----------------------------------------------
*A:7750-1>config>app-assure>group#
```

In the above example, a single leaky-bucket policer is configured with a rate set to 30 kb/s and maximum burst size of 10 kbytes. This provides enough bandwidth to ensure normal operations, while still providing a ceiling limit of how much traffic any eNB can send toward the MMEs.

The value for this policer is a function of the amount of user equipment (UEs) served by the eNB/FAP. For example, in a small cell deployment, with 32 active users per 9962 FAP, the S1-MME bandwidth is estimated to be:

Uplink — toward MME : 2.7 kb/s

Downlink — from MME toward FAP : 28 kb/s

Add the defined policer as a subscriber policy, as follows:

```
A:7750-1>config>app-assure>group>policy>aqp# entry 100
A:7750-1>config>app-assure>group>policy>aqp>entry$ info
----------------------------------------------
                        description "limit SCTP traffic"
                        match
                            traffic-direction subscriber-to-network
                            ip-protocol-num eq sctp
                        exit
                        action
                            bandwidth-policer "sctp_bw_limit"
                            flow-count-limit "sctp_flow_count"
                        exit
                        no shutdown
----------------------------------------------
A:7750-1>config>app-assure>group>policy>aqp>entry$
```

Step 2.4.1 Configure additional limits for all traffic to MMEs.

To further protect the MMEs from a distributed attack, whereby a number of eNBs or FAPs are compromised, an AA FW can be configured to limit total traffic, not just from a single eNB as outlined in previous sections, but from all eNBs toward the MMEs.

It is recommended to configure the following three protection limits:

1. total bandwidth of SCTP toward MMEs
2. total number of flows toward MMEs
3. flow setup rate toward the MMEs

The configuration is shown below:

```
*A:7750-1>config>app-assure>group# info
----------------------------------------------
            policer "limit_total_sctp_bw" type single-bucket-
bandwidth granularity system create
                rate 1200
                mbs 100
            exit
            policer "limit_total_sctp_flows" type flow-count-
limit granularity system create
                flow-count 400
            exit
            policer "limit_total_sctp_flows_rate" type flow-rate-
limit granularity system create
                rate 100
                mbs 100
            exit
----------------------------------------------
*A:7750-1>config>app-assure>group#
```

Note:

- The policers are of type **system** and not **subscriber** in order to be applied to all eNBs or FAPs, as is the case when auto-transit subscribers are created (see Step 1).
- The actual limits of these policers are a function of the total number of eNBs served by the SeGW. In the above configuration, it is assumed that there are 400 eNBs. Therefore, the total limit is 400 flows of SCTP traffic.
- A flow setup rate limit of 100 is set to protect MMEs from a storm of new SCTP flows.

The policers are then referenced from within the appropriate AQP entry that matches the MMEs traffic and SCTP:

```
*A:7750-1>config>app-assure>group>policy>aqp# info
----------------------------------------------
 ---snipped---
                    entry 110 create
                        description " limit system traffic towards MMEs"
                        match
                            traffic-direction subscriber-to-network
                            src-ip eq ip-prefix-list "ALL_eNBs"
```

```
                            dst-ip eq ip-prefix-list "MMEs"
                        exit
                        action
                            bandwidth-policer "limit_total_sctp_bw"
                            flow-rate-limit "limit_total_sctp_flows_rate"
                            flow-count-limit "limit_total_sctp_flows"
                        exit
                        no shutdown
                    exit
*A:7750-1>config>app-assure>group>policy>aqp#
```

Note: It is possible, but redundant, to add the **ip-protocol eq sctp** command as a match condition, because the configured session filter already ensures that only SCTP traffic can flow between eNBs and MMEs.

Step 2.5. Allow only specified SCTP PPIDs toward the MMEs.

In this step, the operator blocks all except the specified SCTP messages that contain configured PPIDs, using an AA SCTP filter configuration:

```
*A:7750-1>config>app-assure>group# sctp-filter
  - no sctp-filter <sctp-filter-name>
  - sctp-filter <sctp-filter-name> [create]

 <sctp-filter-name>    : [32 chars max]
 <create>              : keyword

 [no] description      - Configure a description of the SCTP filter
 [no] event-log        - Configure an event log for packets dropped by the SCTP
                         filter
      ppid             + Configure actions for specific or default PPIDs
                         (Payload Protocol Identifiers)
 [no] ppid-range       - Configure the range of allowable PPIDs for the SCTP
                         filter
```

The filter specifies either a range of PPIDs or individual PPIDs:

```
*A:7750-1>config>app-assure>group>sctp-fltr>ppid$ entry 1
  - entry <entry-id> value <ppid-value> action {permit|deny}
  - no entry <entry-id>

 <entry-id>            : [1..255]
 <ppid-value>          : [0..4294967295]D | [256 chars max]
 <permit|deny>         : permit|deny
```

The PPIDs can be specified either by their values or by names. Names are specified in RFC4960. See Table 6.

Table 6 **SCTP PPIDs**

Value	SCTP PPID	Value	SCTP PPID
0	Reserved by SCTP	31	Service Area Broadcast Protocol (SABP)

Table 6 **SCTP PPIDs (Continued)**

Value	SCTP PPID	Value	SCTP PPID
1	IUA	32	Fractal Generator Protocol (FGP)
2	M2UA	33	Ping Pong Protocol (PPP)
3	M3UA	34	CalcApp Protocol (CALCAPP)
4	SUA	35	Scripting Service Protocol (SSP)
5	M2PA	36	NetPerfMeter Protocol Control Channel (NPMP-CONTROL)
6	V5UA	37	NetPerfMeter Protocol Data Channel (NPMP-DATA)
7	H.248	38	Echo (ECHO)
8	BICC/Q.2150.3	39	Discard (DISCARD)
9	TALI	40	Daytime (DAYTIME)
10	DUA	41	Character Generator (CHARGEN)
11	ASAP	42	3GPP RNA
12	ENRP	43	3GPP M2AP
13	H.323	44	3GPP M3AP
14	Q.IPC/Q.2150.3	45	SSH over SCTP
15	SIMCO <draft-kiesel-midcom-simco-sctp-00.txt>	46	Diameter in a SCTP DATA chunk
16	DDP Segment Chunk	47	Diameter in a DTLS/SCTP DATA chunk
17	DDP Stream Session Control	48	R14P. BER Encoded ASN.1 over SCTP
18	S1 Application Protocol (S1AP)	49	Unassigned
19	RUA	50	WebRTC DCEP
20	HNBAP	51	WebRTC String
21	ForCES-HP	52	WebRTC Binary Partial (deprecated)
22	ForCES-MP	53	WebRTC Binary
23	ForCES-LP	54	WebRTC String Partial (deprecated)
24	SBc-AP	55	3GPP PUA
25	NBAP	56	WebRTC String Empty
26	Unassigned	57	WebRTC Binary Empty

Table 6 **SCTP PPIDs (Continued)**

Value	SCTP PPID	Value	SCTP PPID
27	X2AP	58-4294967295	Unassigned
28	IRCP - Inter Router Capability Protocol		
29	LCS-AP		
30	MPICH2		

It is recommended to limit the SCTP traffic to only those packets with S1 AP PPID. The SCTP filter can be configured to deny all by default and only allow PPID S1 AP (by value = 18 or by name: **s1-application-protocol**) as follows:

```
*A:7750-1# configure application-assurance group 1:1
            sctp-filter "SCTP-PPID-Filter" create
                description "Allow only S1AP PPID"
                event-log "FW_drops_log"
                ppid
                    default-action deny
                    entry 1 value "s1-application-protocol" action permit
                exit
```

This configured SCTP filter is then referenced as an action from within an AQP entry:

```
*A:7750-1>config>app-assure>group>policy>aqp# info
----------------------------------------------
                    entry 100 create
                        description "limit SCTP traffic"
                        match
                            traffic-direction subscriber-to-network
                            ip-protocol-num eq sctp
                        exit
                        action
                            bandwidth-policer "sctp_bw_limit"
                            flow-count-limit "sctp_flow_count"
                            sctp-filter "SCTP-PPID-Filter"
                        exit
                        no shutdown
                    exit
----------------------------------------------
A:7750-1>config>app-assure>group>policy>aqp#
```

To view the packets allowed or denied by the configured SCTP filter:

```
*A:7750-1# show application-assurance group 1:1 sctp-filter "SCTP-PPID-Filter"
===============================================================================
Application Assurance Group 1:1 SCTP Filter "SCTP-PPID-Filter"
===============================================================================
Description           : Allow only S1AP PPID
Maximum PPID          : 4294967295
```

```
Minimum PPID          : 0
Default action        : deny
Configured PPIDs      : 1

Packets arrived       : 0
Packets denied
  Malformed packet    : 0
  PPID out of range   : 0
  PPID denied         : 0
Packets permitted     : 0
===============================================================================
*A:7750-1#
```

Note: The SCTP malformed packet counter shown above increments when an AA SCTP filter encounters an SCTP packet that is malformed, such as:

- IP packet is too small to contain a common SCTP header
- SCTP chunk LEN < 4 bytes: each SCTP chunk header is 4 bytes, so the SCTP chunk cannot be smaller than this
- remaining space in the IP packet is too small to contain a chunk header (for example, your packet has 2 chunks and the 2nd chunk length goes beyond the IP length advertised)
- IP packet is too small to contain the chunk

Currently, the SCTP filter statistics cannot be reset on the fly without shutting down the SCTP filter.

Another way to view the effect of the configured SCTP filter is to check the firewall log, if configured:

```
*A:7750-1# tools dump application-assurance group 1:1 event-
log "FW_drops_log" isa 1/2
```

Step 3. ConfigureAA-ISA to protect SGW (GTP-U traffic).

The steps to configure the AA-ISA in an SeGW to protect against attacks toward the SGW are similar to the steps for SCTP traffic. While GTP filtering is very different from SCTP filtering, configuration to limit the flow counts, bandwidth, and session filter are similar.

Step 3.1. Create an AA IP list for SGWs.

In addition to the lists configured in step 2.1, the operator can optionally configure a list that contains the SGW IP addresses that are served by the SeGW, in case there is more than one.

```
*A:7750-1# configure application-assurance group 1:1 ip-prefix-list "SGWs" create
*A:7750-1>config>app-assure>group>pfx>$ description "Serving Gateways IPs"
*A:7750-1>config>app-assure>group>pfx>$ prefix 172.16.111.1/32
*A:7750-1>config>app-assure>group>pfx>$ prefix 172.16.111.2/32
*A:7750-1>config>app-assure>group>pfx>$ exit
```

Step 3.2. Allow only GTP-U traffic toward SGWs — No port scanning.

Similar to Step 2.2, create an GTP filter to allow only GTP traffic to/from eNBs to SGWs:

```
*A:7750-1>config>app-assure>group>sess-fltr# info
----------------------------------------------
                default-action deny event-log "FW_drops_log"
                ---snipped---
                entry 2 create
                    description "allow GTP-u to SGWs"
                    match
                        ip-protocol-num udp
                        src-ip ip-prefix-list "ALL_eNBs"
                        dst-ip ip-prefix-list "SGWs"
                        dst-port eq 2152
                    exit
                    action permit
                exit
----------------------------------------------
*A:7750-1>config>app-assure>group>sess-fltr#
```

The following session filter needs to be added to the default sub-policy AQP, similar to Step 2.2:

```
*A:7750-1>config>app-assure>group>policy>aqp# info
----------------------------------------------
                    entry 500 create
                        description "apply SeGW session filter rules"
                        match
                            traffic-direction subscriber-to-network
                        exit
                        action
                            session-filter "SeGW_FW"
                            overload-drop
                            error-drop
                            fragment-drop all
                        exit
                        no shutdown
                    exit
                exit
----------------------------------------------
```

For AA to recognize GTP traffic and perform sanity packet checking, configure a GTP filter at the group:partition level:

```
*A:7750-1# configure application-assurance group 1:1 gtp no shutdown
```

Step 3.3. DoS protection — Limit the number of GTP-U flows from eNBs.

AA can be configured to limit the number of GTP flows from an eNB. A GTP-U flow is defined by GTP-U packet destination IP + tunnel ID (TEID).

AA allows the operator to configure two limits: one that applies to the each eNB and one that applies for all GTP-U traffic from all eNBs:

```
*A:7750-1>config>app-assure>group# info
----------------------------------------------
            policer "GTPu-Flow-count-limit" type flow-count-
limit granularity subscriber create
```

```
                flow-count 800
                gtp-traffic
            exit
```

The actual value of the flow count limit is a function of the number of UEs or devices served by an eNB or FAP. In the above case, it is assumed that there are 100 devices with a maximum of 8 GTP-U flows per device. For FAP, the number is typically around 32 devices per FAP. Note: By 3GPP standards, the maximum number of GTP-U tunnels per device is 16.

Assuming that there are 1000 eNBs or FAPs that are served by the SeGW, then to limit the total number of GTP-U flows, the operator can apply the following system policer:

```
*A:7750-1>config>app-assure>group# info
----------------------------------------------
            policer "limit_total_GTPU_Flow_count" type flow-count-
limit granularity system create
                flow-count 800000
                gtp-traffic
            exit
```

Configure AQPs to execute the policers:

```
*A:7750-1>config>app-assure>group>policy>aqp# info
----------------------------------------------
---snipped---
                    entry 120 create
                        description "limit GTP-U traffic"
                        match
                            traffic-direction subscriber-to-network
                        exit
                        action
                            flow-count-limit "GTPu-Flow-count-limit"
                        exit
                        no shutdown
                    exit
                    entry 130 create
                        description "limit TOTAL GTPU towards SGWs"
                        match
                            traffic-direction subscriber-to-network
                        exit
                        action
                            flow-count-limit "limit_total_GTPU_Flow_count"
                        exit
                        no shutdown
                    exit
```

For GTP-U flow count policing, it is important that **aqp-initial-lockup** is enabled:

```
*A:7750-1# configure application-assurance group 1:1 aqp-initial-lookup
```

The above configured limits are applied only to upstream traffic, to protect the network. No limit is placed on the downstream traffic toward the eNBs.

Note: Note: For small cell deployments, the number of GTP-U tunnels per FAP is a function of:

1. deployment mode:
   a. residential = 32 (9962 MSEC-MS-MCI Enterprise) UEs,
   b. enterprise = 8 (9961 MSHC) UEs.
2. number of guaranteed bit rate (GBR) tunnels (max 8) and non-GBR tunnels (max 8) per UE.

Therefore, the GTP-U tunnel limit per FAP should be set to 32 x 8 = 256 for residential deployments or 8 x 8 = 64 for enterprise deployments.

The operator can view the effect of the configured policers on GTP traffic by running the following show routine:

```
*A:7750-1>show>app-assure>group# gtp
===============================================================================
Application Assurance Group 1:1 GTP
===============================================================================
Admin status : Up
Event log    : (Not Specified)
-------------------------------------------------------------------------------
GTP Statistics                              sub-to-net             net-to-sub
-------------------------------------------------------------------------------
Incoming packets                                     0                      0
Packets denied
  UDP packet length                                  0                      0
  GTP message length                                 0                      0
  GTP version                                        0                      0
-------------------------------------------------------------------------------
Packets permitted                                    0                      0
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
GTP Policing Statistics                     sub-to-net             net-to-sub
-------------------------------------------------------------------------------
Packets arrived                                      0                      0
Packets denied
  gtp-traffic flow-count policer                     0                      0
  Other                                              0                      0
-------------------------------------------------------------------------------
Packets permitted                                    0                      0
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
GTP Filter Statistics                       sub-to-net             net-to-sub
-------------------------------------------------------------------------------
Packets arrived                                      0                      0
Packets denied (gtp-filter)                          0                      0
Packets permitted
  gtp-filter                                         0                      0
  no gtp-filter                                      0                      0
-------------------------------------------------------------------------------
Total GTP packets permitted                          0                      0
===============================================================================
*A:7750-1>show>app-assure>group#
```

In the last section shown above, GTP filter statistics are related to GTP filters that are discussed and configured later in Step 3.5 of this chapter.

Step 3.4. DoS protection: Limit the GTP-U bandwidth from eNBs.

This step is similar to Step 3.3, but instead of configuring a flow count policer, the operator configures bandwidth policers:

```
*A:7750-1>config>app-assure>group# info
----------------------------------------------
            policer "GTPU_bw_limit" type single-bucket-
bandwidth granularity subscriber create
                rate 5000
                mbs 100
            exit

            policer "limit_total_GTPU_bw" type single-bucket-
bandwidth granularity system create
                rate 2000000
                mbs 2000
            exit

*A:7750-1>config>app-assure>group>policy>aqp# info
----------------------------------------------
---snipped---
                    entry 120 create
                        description "limit GTP-U traffic"
                        match
                            traffic-direction subscriber-to-network
                        exit
                        action
                            bandwidth-policer "GTPU_bw_limit"
                            flow-count-limit "GTPu-Flow-count-limit"
                        exit
                        no shutdown
                    exit
                    entry 130 create
                        description "limit TOTAL GTPU towards SGWs"
                        match
                            traffic-direction subscriber-to-network
                        exit
                        action
                            bandwidth-policer "limit_total_GTPU_bw"
                            flow-count-limit "limit_total_GTPU_Flow_count"
                        exit
                        no shutdown
                    exit
```

The above configured limits are applied only to upstream traffic, to protect the network. No limit is placed on downstream traffic toward the eNB.

As a debugging tool, the operator can use the AA **flow-record-search** command to check the status of GTP flows through the ISA:

```
*A:7750-1# tools  dump application-assurance group 1:1 flow-record-search  isa 1/
2 flow-status active protocol "gtp"
=====================================================
Application-Assurance flow record search, Version 1.0
Search Start Time:     "06/16/2015 20:38:09" (UTC)
```

```
 Search Criteria:
  group[:partition]:   1:1
  isa:                 1/2
  protocol name:       "gtp"
  application name:    none specified
  app-group name:      none specified
  flow-status:         active
  start-flowId:        none specified
  classified:          none specified
  server-ip:           none specified
  server-port:         none specified
  client-ip:           none specified
  bytes-tx:            none specified
  flow-duration:       none specified
  max-count:           none specified
  search-type:         default
=====================================================
FlowId  Init  Src-ip                                    Dst-
ip                                 Ip-prot     Src-prt Dst-
prt Protocol               Application            Pkts-tx    Bytes-
tx             Pkts-disc  Bytes-disc Time-ofp(UTC)         Time-olp(UTC)
SEARCH COMPLETED.
Search End Time:     "06/16/2015 20:38:09" (UTC)
Total Records:       0
=====================================================
*A:7750-1#
```

GTP flows that are to be denied by the previous AA configurations should not appear in the search results.

Step 3.5. Further GTP filtering and validation.

AA allows the operator to configure a GTP filter to enforce which GTP message types are allowed/denied, as well as the maximum allowed GTP message length:

```
*A:7750-1>config>app-assure>group>gtp>gtp-fltr#
 [no] description     - Configure a description of the GTP filter
 [no] event-log       - Configure an event log for packets dropped by the GTP filter
 [no] max-payload-le* - Configure the maximum payload length of the GTP filter
      message-type    + Configure actions for specific or default messages

*A:7750-1>config>app-assure>group>gtp>gtp-fltr#
```

Note: An AA GTP filter allows the operator to configure a maximum payload size for the GTP traffic. However, in this configuration example, no maximum payload size is configured.

The list of GTP message types are defined by 3GPP standard 3GPP TS 29.281 as per Table 7.

Table 7 **GTP Messages**

Message Type Value (Decimal)	Message	Message Type Value (Decimal)	Message
1	“echo-request”	55	“”forward-relocation-complete
2	“echo-response”	56	“relocation-cancel-request”
3	“versi”n-not-supported:	57	“relocation-cancel-response”
4	“node-alive-request”	58	“forward-sms-context”
5	“ node-alive-response”	59	“forward-relocation-complete-acknowledge”
6	“redirection-request”	60	“forward-sms-context-acknowledge”
7	“redirection-response”	70	“ran-information-relay”
16	“create-pdp-context-request”	96	“mbms-notification-request”
17	“create-pdp-context-response”	97	“mbms-notification-response”
18	“update-pdp-context-request”	98	“mbms-notification-reject-request”
19	“update-pdp-context-response”	99	“mbms-notification-reject-response”
20	“delete-pdp-context-request”	100	“create-mbms-context-request”
21	“delete-pdp-context-response”	101	“create-mbms-context-response”
22	“initiate-pdp-context-activation-request”	102	“update-mbms-context-request”
23	“initiate-pdp-context-activation-response”	103	“update-mbms-context-response”
26	“error-indication”	104	“delete-mbms-context-request”
27	“pdu-notification-request”	105	“delete-mbms-context-response”
28	“pdu-notification-response”	112	“mbms-registration-request”
29	“pdu-notification-reject-request”	113	“mbms-registration-response”
30	“pdu-notification-reject-response”	114	“mbms-de-registration-request”
31	“supported-extension-headers-notification”	115	“mbms-de-registration-response”
32	“send-routing-information-for-gprs-request”	116	“mbms-session-start-request”

Table 7 **GTP Messages (Continued)**

Message Type Value (Decimal)	Message	Message Type Value (Decimal)	Message
33	“send-routing-information-for-gprs-response”	117	“mbms-session-start-response”
34	“Failure-report-request”	118	“mbms-session-stop-request”
35	“failure-report-request”	119	“mbms-session-stop-response”
36	“note-ms-gprs-present-request”	120	“mbms-session-update-request”
37	“note-ms-gprs-present-response”	121	“mbms-session-update-response”
48	“identification-request”	128	“ms-info-change-notification-request”
49	“identification-response”	129	“ms-info-change-notification-response”
50	“sgsn-context-response”	240	“data-record-transfer-request”
51	“sgsn-context-request”	241	“data-record-transfer-response”
52	“sgsn-context-acknowledge”	254	“end-marker”
53	“forward-relocation-request”	255	“g-pdu”
54	“forward-relocation-response”		

Of the 67 GTP message types shown above, only 6 are allowed, by the standards, for GTP-U:

```
echo-request      echo-response      error-indication
g-pdu             end-marker         supported-extension-headers-notification
```

If these message types are permitted by the configured GTP filter, AA performs extensive GTP-U header checking on these six types.

Note: If no GTP filter is configured, no extensive GTP-U header checks are performed. For example, if the operator wants to allow all GTP-U packets and perform all GTP header sanity checks, then a GTP filter that permits all message types needs to be configured, with the default action of permit and with no values, such as:

```
gtp-filter “allow-all” create
     message-type
        default-action permit
```

Because AA FW in an SeGW is protecting an S1-U interface running GTP-U, the GTP filter only needs to allow the six GTP messages that are permitted for GTP-U:

```
*A:7750-1>config>app-assure>group>gtp# info
----------------------------------------------
                gtp-filter "filter-gtp-msgs" create
                    description "allow only certain msg types"
                    message-type
                        default-action deny
                        entry 1 value "echo-request" action permit
                        entry 2 value "echo-response" action permit
                        entry 3 value "error-indication" action permit
                        entry 4 value "supported-extension-headers-
notification" action permit
                        entry 5 value "end-marker" action permit
                        entry 6 value "g-pdu" action permit
                    exit
                exit
                no shutdown
----------------------------------------------
*A:7750-1>config>app-assure>group>gtp#
```

This GTP filter is then referenced from within an AQP entry action, as follows, in order for it to take effect:

```
*A:7750-1>config>app-assure>group>policy>aqp# info
----------------------------------------------
                    entry 120 create
                        description "limit GTP-U traffic"
                        match
                            traffic-direction subscriber-to-network
                            dst-ip eq ip-prefix-list "SGWs"
                        exit
                        action
                            bandwidth-policer "GTPU_bw_limit"
                            flow-count-limit "GTPu-Flow-count-limit"
                            gtp-filter "filter-gtp-msgs"
                        exit
                        no shutdown
                    exit
```

The operator can view the effect of the configured GTP filter on S1-U traffic using the following show routine:

```
*A:7750-1>show>app-assure>group# gtp gtp-filter "filter-gtp-msgs"
===============================================================================
Application Assurance Group 1:1 GTP Filter "filter-gtp-msgs"
===============================================================================
Description            : allow only certain msg types
Maximum payload length : (Not Specified)
Default action         : deny
Configured messages    : 6

Packets arrived        : 0
Packets denied
  Payload length       : 0
  Message type         : 0
```

```
  Mandatory header     : 0
  Extension header     : 0
  Information element  : 0
Packets permitted      : 0
===============================================================================
*A:7750-1>show>app-assure>group#
```

The above output is in addition to the information provided by the overall GTP show command:

```
*A:7750-1>show>app-assure>group# gtp
```

Step 4. Configure AA-ISA to protect NMS (OAM Traffic).

Step 4.1. Create an IP AA list that contains the NMS server IPs.

```
*A:7750-1# configure application-assurance group 1:1
*A:7750-1>config>app-assure>group# info
----------------------------------------------
            ip-prefix-list "NMSs" create
                description "Network Management-OAM subnet"
                prefix 172.16.120.0/30
            exit
```

Note: In the case of small cell deployments, different NMS servers need to be configured.

Step 4.2. Allow eNBs to initiate FTP- and ICMP-only traffic toward NMS, block port scanning.

```
*A:7750-1>config>app-assure>group>sess-fltr# info
----------------------------------------------
                default-action deny event-log "FW_drops_log"

                entry 3 create
                    description "allow FTP to NMS"
                    match
                        ip-protocol-num tcp
                        src-ip ip-prefix-list "ALL_eNBs"
                        dst-ip ip-prefix-list "NMSs"
                        dst-port eq 22
                    exit
                    action permit
                exit
                entry 4 create
                    description "allow ICMP to NMS"
                    match
                        ip-protocol-num icmp
                        src-ip ip-prefix-list "ALL_eNBs"
                        dst-ip ip-prefix-list "NMSs"
                    exit
                    action permit
                exit
----------------------------------------------
*A:7750-1>config>app-assure>group>sess-fltr#
```

The operator can view the effect of the session filter on traffic, in terms of how many times it is applied, using the following show routine:

```
*A:7750-1>show>app-assure>group# session-filter
===============================================================================
AA Session Filter Table
===============================================================================
Name                            Default Action    Referenced           Entries
-------------------------------------------------------------------------------
SeGW_FW                         deny              aqp                        4
-------------------------------------------------------------------------------
No. of session filters: 1
===============================================================================
*A:7750-1>show>app-assure>group#
*A:7750-1>show>app-assure>group# session-filter "SeGW_FW"
===============================================================================
AA Session Filter Instance "SeGW_FW"
===============================================================================
Description    : (Not Specified)
Default Action : deny
    Event Log  : FW_drops_log
AQP Entries    :
          500
-------------------------------------------------------------------------------
Filter Match Criteria
-------------------------------------------------------------------------------
Entry          : 1
Description    : allow SCTP to MMEs
IP Protocol    : sctp
Source IP List : ALL_eNBs
Dest IP List   : MMEs
Action         : permit
    Event Log  : (Not Specified)
Hits           : 0 flows
-------------------------------------------------------------------------------
Entry          : 2
Description    : allow GTP-u to SGWs
IP Protocol    : udp
Source IP List : ALL_eNBs
Dest IP List   : SGWs
Dest Port      : eq 2152
Action         : permit
    Event Log  : (Not Specified)
Hits           : 0 flows
-------------------------------------------------------------------------------
Entry          : 3
Description    : allow FTP to NMS
IP Protocol    : tcp
Source IP List : ALL_eNBs
Dest IP List   : NMSs
Dest Port      : eq 22
Action         : permit
    Event Log  : (Not Specified)
Hits           : 0 flows
-------------------------------------------------------------------------------
Entry          : 4
Description    : allow ICMP to NMS
IP Protocol    : icmp
Source IP List : ALL_eNBs
```

```
Dest IP List    : NMSs
Action          : permit
    Event Log   : (Not Specified)
Hits            : 0 flows
-------------------------------------------------------------------------------
No. of entries    : 4
===============================================================================
*A:7750-1>show>app-assure>group#
```

Note: The above configuration is generic and may need to be modified to suit the deployment requirements. For example, in the case of small cell SeGW deployment, traffic on other ports needs to be allowed to/from different NMS type servers, such as allowing TCP port 7003 and port 7013 to HDM servers. This can be accomplished by configuring additional entries in the above session filter.

Note: By allowing port 22 for FTP, the AA FW automatically opens and closes the associated data channel ports. For more information about AA FW capabilities, with regard to OAM FW protection, see Application Assurance Stateful Firewall.

Conclusion

The SR OS AA stateful firewall feature runs on AA-ISA and extends application-level analysis to provide an in-line stateful service, integrated within the Security Gateway (SeGW).

AA stateful packet filtering, combined with AA Layer 7 classification and control, provides advanced, next-generation firewall functionality, protecting mobile network core infrastructure, such as MMEs, SGWs, and NMSs.

Application Assurance — Stateful Firewall

This chapter describes Application Assurance stateful firewall (FW) configurations for protecting residential and WiFi subscribers.

Topics in this chapter include:

Applicability

Initially, this chapter was written for SR OS release 11.0.R1. The TCP validation section was added for SR OS release 14.0.R4.

Overview

The AA SR OS 11.0.R1 stateful FW feature extends AA-ISA application level analysis to provide an in-line integrated stateful service that protects subscribers from malicious attacks. AA stateful packet filtering combined with AA L7 classification and control, empowers operators with advanced, next generation firewall functionality that is integrated within the Service Router. The AA stateful firewall (FW) and application firewall runs on AA-ISA. Using stateful inspection, the AA firewall not only inspects packets at layers 3-7, but also monitors and keeps track of the connection's state. If the operator configures a **deny action** within a session filter, then the matching packets (matching both the AA Application QoS policy (AQP) and associated session filter match conditions) are dropped and no flow session state/ context is created.

AA FW can be used in all deployments of AA-ISA; mobile (MG OS) and fixed (SR OS), however the configurations examples here, while still very applicable (and almost 100% identical in mobile deployments) are focused on AA-ISA deployments in fixed networks.

The AA-ISA FW enabled solution provides:

- Stateful (and stateless) packet filtering and inspection with application-level gateway (ALG) support
- DoS attack protection

In SR OS release 14.0, additional firewall functionalities were added, such as TCP-validation, threshold crossing alerts, syslog and statistics related to firewall actions.

The objective of this chapter is to describe the required configuration within AA-ISA (divert to AA-ISA basic knowledge is assumed) in order to enable AA FW and protect AA subscribers from attacks (Unsolicited attacks and DoS attacks), while still allowing pin-holing through the firewall, so that applications like peer to peer gaming and various ALGs (such as FTP) are not affected.

Stateful Filtering

By performing stateful inspection, AA-ISA takes into account which side initiated a session and acts accordingly. Stateful flow processing and inspection utilizes IP layers 3/ 4 header information to build a state of the flow within AA-ISA. Layer 7 inspection is used in order to provide ALG support. Stateful flow/session processing takes note of the originator of the session and hence can allow traffic to be initiated from the subscriber, while denying (when configured) traffic originating from the network. Packets received from the network are inspected against the session filter and only those that are part of a subscriber initiated session are allowed.

Figure 22 **Block Unsolicited Traffic**

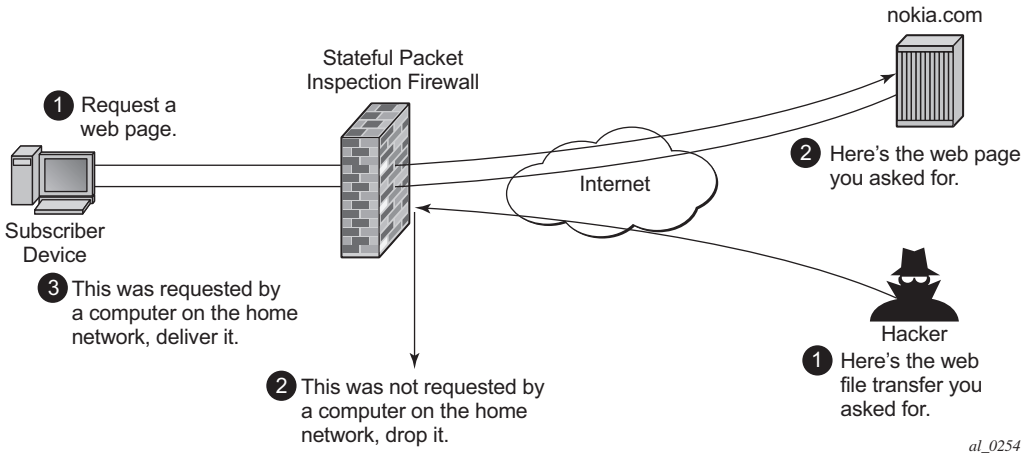

To support the example shown in Figure 22, AA is configured with an action to block unsolicited traffic; traffic that is not originated/initiated from the subscriber. The direction field in match criteria of AQPs is utilized to enable this functionality.

Figure 23 **SFW — Allow Gaming**

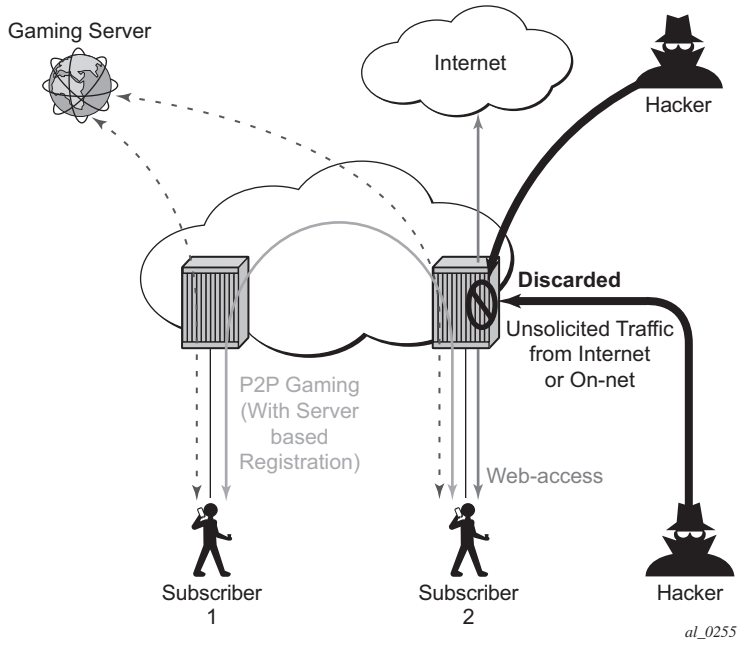

Figure 23 shows a similar concept. It is used to allow UDP traffic for peer to peer applications (such as gaming). Once the traffic from one peer is seen by AA-ISA, a pin-hole is opened in the reverse direction to allow for the corresponding UDP traffic from the peer.

Stateless packet filtering on the other hand does not take note of the session initiator. It discards or allows packets independently of any previous packets. In addition to AA-ISA's support for stateless (and stateful) filtering, stateless packet filtering can be performed in the system using line card ACLs (and/or MGISM PCC rules in mobile gateway deployments).

Application Layer Gateway Filtering

Figure 24 **ALG Support Example — FTP**

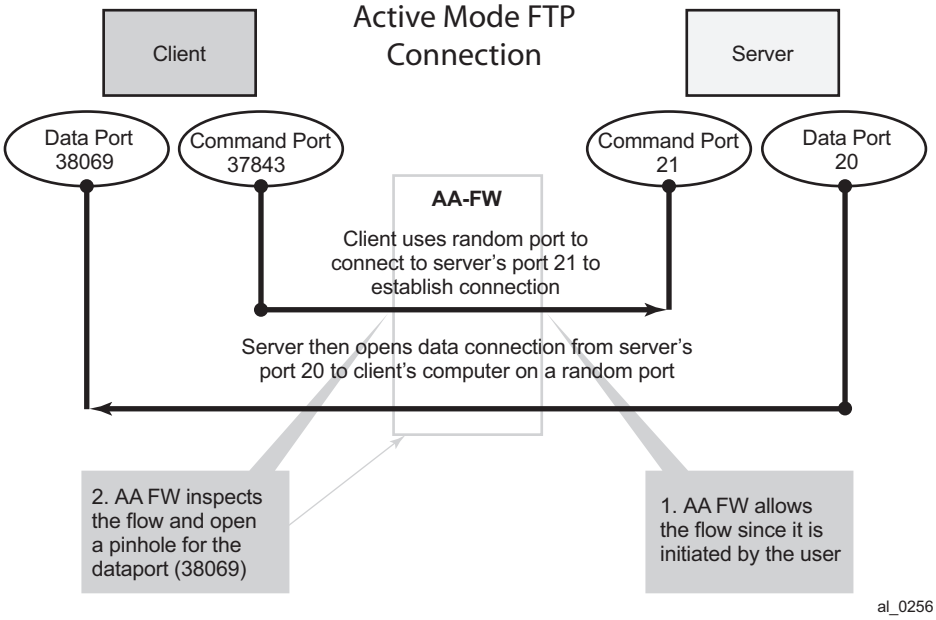

AA FW inspection of packets at Layer 7 offers Application Layer Gateway functionality for a large list of applications (for example, FTP, SIP, RTSP, PPTP, IRC, etc.). These applications make use of control channels or flows that spawn other flows. AA FW inspects the payload of these control flows so it can open a pinhole in advance for unspawned data flows. Figure 24 depicts an example of AA ALG support for FTP traffic.

Denial Of Service (DOS) Protection

DoS attacks work by consuming network and system resources, making them unavailable for legitimate network applications. Network flooding attacks, malformed packets and port scans are examples of such DoS attacks.

The aim of AA FW DOS protection is to protect subscribers and prevent any abuse of network resources.

Using AA FW stateful session filters, operators can protect their subscribers from any port scan scheme. This can be done by configuring the session filters to disallow any traffic that is initiated from the network.

Furthermore, AA ISA provides configurable flow policers. These policers, once configured, prevent a wide range of flooding attacks (such as ICMP PING flooding, UDP flooding, SYN Flood Attack...etc.). These policers provide protection at multiple levels; per system per application/application groups and per subscriber per applications/applications groups.

There are two types of AA ISA flow policers; flow setup rate policers and flow count policers. Flow setup rate policers limit the number of new flows, while flow count policers limit the total number of active flows.

In order to protect hosts and network resources, AA FW validates/checks different fields in the packet's header (checksum, TCP Flag, etc.) and if any fails it declares the packet to be invalid. This complements the 7x50 subscriber management enhanced security features, such as IP (or MAC) anti-spoofing protection (such as protecting against LAND attacks) and network protocol DoS protections. The cut-through-drop AQP action must be configured in order to drop these types of invalid packets.

Virtual FW/Zone-Based FW

AA FW can provide up to 128 virtual FWs, each with its own FW policies. This is achieved through the use of AA-partitions.

In addition, AA subscribers within the same AA partition can have different application profiles with different Application Service Options (ASO) values. This provides a further control mechanism to enable/disable firewall rules.

For example, the operator may want to have some subscribers possess full firewall protection, while other subscribers not subscribed to this service to have a partial firewall protection that focuses on protecting network resources, rather than network and subscribers resources.

Configuration

AA-ISA AQPs were enhanced in R11.0.R1 with several AQP actions that provide session filtering functionality. As is the case of all AQPs, these have partition level scope, which allows different FW policies to be implemented by utilizing AA partitions concepts within the same AA-ISA group. Hence, multiple virtual AA FW instances can be realized, without the need for multiple physical instances of FWs to implement different FW policies.

The AA FW stateful session filter consists of multiple entries (similar to ACLs) with a match an action per entry. Actions are **deny** or **permit**. A **deny** action results in packets being discarded without creating a session/flow context. Match conditions include IP protocol types, source and destination IP addresses and ports. An overall default action is also configurable in case of no match to any session filter entry.

AQPs with session filter actions need to have — as a matching condition — traffic direction, ASOs, and/or a subscriber name. These AQP match rules cannot have any references to applications and/or application groups.

An AQP action to drop malformed/errored packets is also available.

Statistics are incremented when packets are dropped by a session filter. These are accounted against:

- protocol = denied by default policy,
- application= unknown,
- application group = unknown.

The configuration topology is shown in Figure 25.

Figure 25 **Configuration Topology**

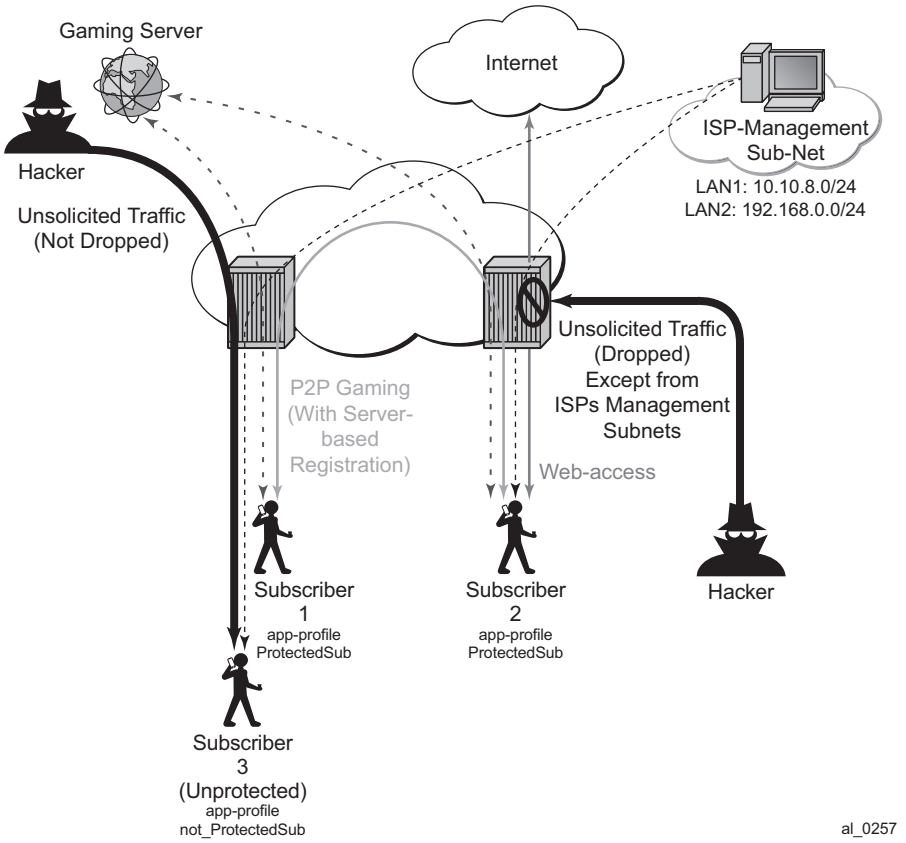

Step 1. Application Profile configuration:

There is nothing new introduced in application profiles in order to support FW. This section explains how to configure the application profile to allow differentiated FW services for different subscribers. In a nut shell, the AA common building construct/attribute for differentiated policy is ASO.

To configure an ASO for FW protection:

```
configure application-assurance group 1:1 policy
  begin
  app-service-options
     characteristic "FW-Protection" create
            value "None"
            value "ON"
            default-value "None"
     exit
     characteristic "ISP-Protection" create
            value "None"
            value "ON"
            default-value "None"
     exit
     characteristic "DOS-Protection" create
            value "None"
```

```
            value "ON"
            default-value "None"
        exit
    exit
```

In the preceding example:

- ASO FW protection allows the operator to select if the subscriber is FW protected or not.
- ASO DOS protection refers to if the subscriber is protected from DOS attacks.
- ASO ISP protection is different from the preceding two as it protects the ISP resources by (in the example that follows) not allowing unsolicited traffic. This should be ON for all subscribers (it is then arguable if someone needs it to be defined in the ASO list, instead of merely configuring an AQP to protect ISP resources all the time).

These ASOs are referenced in appProfiles (and later in AQPs) as follows:

```
configure application-assurance group 1:1 policy
    begin
        app-profile "Protected" create
           divert
           characteristic "FW-Protection" value "ON"
           characteristic "ISP-Protection" value "ON"
           characteristic "DOS-Protection" value "ON"
    exit
```

The preceding application profile Protected is assigned to subscribers who opted/subscribed to the firewall protection service; for example sub 1 and sub 2 in the example shown in Configuration Topology.

Subscribers who are not protected (for example sub 3 in Figure 25) are assigned a different profile:

```
configure application-assurance group 1:1 policy
    begin
        app-profile "unProtected" create
           divert
           characteristic "FW-Protection" value "ON"
           characteristic "ISP-Protection" value "ON"
           characteristic "DOS-Protection" value "ON"
    exit
```

An alternative method to using application profiles/ASOs to provide differentiated services is to configure multiple partitions with different AQPs/ session filters. One partition for example will be for subscribers who are provided with firewall protection, while another is used for subscribers who are not protected. This configuration is simpler and provides statistics per partition. This example however covers the more complex case using ASOs/appProfiles.

Step 2. Flow count policer configuration:

```
configure application-assurance group 1 policer Dos_police_Flow_count type flow-
count-limit granularity subscriber create
   flow-count 500
exit
```

The preceding configuration limits the number of flows a subscriber can have at any time to 500. This is done to protect against DoS attacks. The value 500 is arbitrary and requires tuning for each deployment.

```
configure application-assurance group 1 policer Dos_Police_ICMPFlows type flow-
count-limit granularity system create
   flow-count 5000
exit
```

This configuration limits the total number of flows that matches the configured AQP matching condition. It is used for ICMP applications to prevent mass port scanning.

Step 3. TCP Protocol Validation configuration

```
configure application-assurance group 1:1 tcp-validate TCP_protect create
```

This simple configuration allows the operator to call TCP_protect policy from within an AQP action entry.

The operator can also configure the policy to be “strict”, in which case the AA checks for valid sequence and acknowledgements numbers. In this example, the “strict” option is not used.

Step 4. Application configuration

The following configuration is standard with AppDB. It is shown here for reference.

```
configure application-assurance group 1:1 policy begin
  application ICMP create
  exit
  app-filter
   entry 1540 create
    protocol eq "non_tcp_udp"
    ip-protocol-num eq icmp
    application "ICMP"
    no shutdown
```

```
exit
entry 35500 create
 protocol eq "non_tcp_udp"
 ip-protocol-num eq ipv6-icmp
 application "ICMP"
 no shutdown
exit
```

Step 5. Session-Filter

The following displays session-filter configuration commands to be used in Step 6 later.

```
configure application-assurance group 1:1 session-filter <name> create
description  <description>
    default-action permit|deny    # default=deny
    entry n create
        description <entry-description>
        match
            ip-protocol-num <ip-protocol-number>
            no src-ip  <ip4_or_v6-address/mask>
            no dst-ip  <ip4_or_v6-address/mask>
            no src-port {eq|gt|lt} <port-num> #or
                    range <start-port-num> <end-port-num>
            no dst-port {eq|gt|lt} <port-num> #or
                    range <start-port-num> <end-port-num>
        exit
        action permit|deny
    exit
    entry m create
    ---snip---
```

Parameters

- **entry** *n* — A session filter can have multiple match-action rules, each of these match-action rules represent an entry within the session-filter. The entries are executed in order. If a match is found, within one entry, the subsequent entries within the session-filter are skipped (not evaluated).
- **default-action** [**permit** | **deny**] — This action is performed if no match is found for any of the configured entries within the session-filter. Default is deny.
  - A **deny** action will drop the packet and will not allow a flow record to be allocated for that flow. A **drop** action within AA AQP will drop the packet but it will still create flow record.
  - A **permit** action will allow the packet to flow through the system. A flow record is also allocated. The packet may get dropped by other configured AQP actions (due to header check failures).
- **description** *description-string*

This configures a text string, up to 80 characters, which can be used to describe the use of the session-filter.

– **match** — Keywords to perform the action specified under the **action** keyword only if the conditions in the match section are met.

• **ip-protocol** *ip-protocol-number*

ip-protocol-number — 1..255

- Decimal, hexadecimal or binary representation

- Supported IANA IP protocol names:

- crtp, crudp, egp, eigrp, encap, ether-ip, gre, icmp, idrp, igmp, igp, ip, ipv6, ipv6-frag, ipv6-icmp, ipv6-no-nxt, ipv6-opts, ipv6-route, isis, iso-ip, l2tp, ospf-igp, pim, pnni, ptp, rdp, rsvp, sctp, stp, tcp, udp, vrrp

• **src-ip**/**dst-ip** *ipv4-address/mask***src-ip**/**dst-ip** *iv6-address/mask*

- Source/destination IP address within the packet header.

- IPv4 or IPv6 formats are allowed, with prefixes masks.

• **src-port** *src-port-numbers*

src-port {**eq** | **gt** | **lt**} *port-num*

eq — equal, exact match

gt — match port numbers that are greater than the one specified.

lt — match port numbers that are smaller than the one specified.

port-num — 0..65535 (Applicable to TCP, UDP and SCTP protocols only.)

• **src-port range** *start-port-num end-port-num*

range — Keyword- that match port numbers within the specified range:

start-port-num — 0..65535

end-port-num — 0..65535

• **dst-port** *dst-port-number*

- Same as source port number explained above, but applied against destination port number.

– **action deny** | **permit**

• **deny** or **permit** action is only executed if a match is found.

• **deny** action will drop the packet and will not create a flow record.

• **permit** action will allow the packet to go through (unless another different action is found that causes it to be dropped).

– **no entry** *entry-id*

• Causes the entry to be deleted.

– **no session-filter** *session-filter-name*

• Causes the session filter to be deleted.

```
config application-assurance group 1:1
```

```
        session-filter " denyUnsolictedwMgntCntrl" create
        description "S-FW opted-in sub – allow ISP access"
        default-action deny
        entry 10 create
            description "allow ICMP access from ISP LAN1"
            match
                ip-protocol-num icmp
                src-ip 10.10.8.0/24
            exit
            action permit
         exit
         entry 20 create
             description "allow ICMP access from ISP LAN2"
            match
                ip-protocol-num icmp
                src-ip 192.168.0.0/24
            exit
            action permit
        exit
        entry 30 create
            description "allow all TCP (e.g. FTP/telnet)access from ISP LAN2"
            match
                ip-protocol-num tcp
                src-ip 192.168.0.0/24
            exit
            action permit
        entry 40 create
            description "allow TCP on port 80 /HTTP access from ISP LAN1"
            match
                ip-protocol-num tcp
                src-ip 10.10.8.0/24
                dst-port eq 80
            exit
            action permit
        exit
```

This session filter is used to protect systems located in LAN2. It drops all unsolicited traffic except for FTP coming from LAN1.

```
configure application-assurance group 1:1
   session-filter "protectISPLan2" create
        description "S-FW to deny all unsolicited requests to LAN2"
        default-action deny
        entry 10 create
            description "allow ftp access from ISP LAN1"
            match
                ip-protocol-num tcp
                src-ip 10.10.8.0/24
                dst-port eq 21
            exit
            action permit
        exit
   exit
```

Step 6. AQP configuration:

```
configure application-assurance group 1:1  policy
  begin
  app-qos-policy

    entry 100 create
        description "Protecting ISP1 from DoS attacks from subs"
        match
            traffic-direction subscriber-to-network
            characteristic "ISP-Protection" eq "ON"
            dst-ip eq 10.10.8.0/24
         exit
         action
            flow-count-limit Dos_police_Flow_count
        exit
        no shutdown
   exit

    entry 105 create
        description "Protecting ISP2 from DoS attacks from subs"
        match
            traffic-direction subscriber-to-network
            characteristic "ISP-Protection" eq "ON"
            dst-ip eq 192.168.0.0/24
        exit
        action
            flow-count-limit Dos_police_Flow_count
        exit
        no shutdown
    exit
```

These AQPs protect the ISP network by limiting the number of concurrent flows. Dropping malformed packets is done by entry 130 (later).

To guard against ICMP flooding attacks, a flow count policer (defined earlier) is used as follows:

```
configure application-assurance group 1:1 policy
  begin
  app-qos-policy entry 107 create
      match
          application eq "ICMP"
          traffic-direction subscriber-to-network
      exit
      action
          flow-count-limit Dos_Police_ICMPFlows
      exit
      no shutdown
    exit
```

To guard against attacks exploiting TCP handshake mechanisms, TCP validate policy (defined earlier) is used as follows:

```
configure application-assurance group 1:1 policy
  begin
    app-qos-policy
        entry 108 create
            match
                characteristic "ISP-Protection" eq "ON"
```

```
            exit
            action
                tcp-validate "TCP_protect"
            exit
            no shutdown
        exit
        entry 109 create
            match
                characteristic "FW-Protection" eq "ON"
            exit
            action
                tcp-validate "TCP_protect"
            exit
           no shutdown
       exit
```

TCP validation works on both direction and needs to be called in from within a sub-default AQP entry. Therefore, this AQP action cannot be restricted to one ISP versus another because no destination IP can be specified. The configuration shown runs TCP validation policy when ISP-Protection or FW-protection ASOs are enabled.

The preceding configuration will ensure, for example, that no TCP session starts without the proper handshake message exchanges.

In order to protect ISP LAN2 from all incoming traffic (unsolicited), the operator configures entry 120.

```
    entry 120 create
        match
            traffic-direction subscriber-to-network
            characteristic "ISP-Protection" eq "ON"
        exit
        action
            session-filter "protectISPLan2"
        exit
        no shutdown
    exit
```

ProtectISPLan2 session filter drops all unsolicited traffic to LAN2 (highly secure) except for access to FTP services coming from ISP LAN1. Details of these configurations are shown in Session-Filter (step 5).

To enable stateful protection for opted-in subscribers:

```
configure application-assurance group 1:1 policy
  begin
    app-qos-policy
        entry 110 create
            description "FW for managed opted-in subs"
            match
                traffic-direction network-to-subscriber
                characteristic "FW-Protection" eq "ON"
            exit
            action
                session-filter "denyUnsolictedwMgntCntrl"
            exit
```

```
            no shutdown
        exit
```

The preceding AQP protects opt-in subscribers from unsolicited traffic but still allows unsolicited traffic from ISP subnets to manage the subscriber's network.

Dropping malformed/illegal packets and protecting against DOS attacks is done via the following entry 130 and 131.

```
    entry 130 create
        match
            traffic-direction subscriber-to-network

            characteristic "DOS-Protection" eq "ON"
        exit
        action
            flow-count-limit Dos_police_Flow_count
        exit
        no shutdown
    exit
    entry 131 create
        match
            characteristic "DOS-Protection" eq "ON"
        exit
        action
            error-drop
            overload-drop
            fragment-drop all
        exit
        no shutdown
    exit
```

Step 7. Configuration of Threshold Crossing Alerts (TCA).

Operators can configure AA to generate TCAs for various firewall related parameters, such as error-drop, session-filter hits, TCP-validate, fragment-drop-all etc. as well as flow count policers. An example of a TCA used for TCP_validation policy is as follows:

```
configure application-assurance group 1:1 statistics  threshold-crossing-alert
        tcp-validate "TCP_protect" direction from-sub create
        high-wmark 50 low-wmark 40
    exit
```

Unlike the other TCAs, in order to configure TCAs for flow count policers, operators need first to configure AA admit-deny to allocate ISA resources to record, such as:

```
configure application-assurance group 1:1 statistics aa-admit-deny policer-stats-
resources
```

Then, a TCA can be configured for any flow based policer in the system, such as:

```
configure application-assurance group 1:1 statistics  threshold-crossing-alert
```

```
        policer "Dos_police_Flow_count" direction from-sub create
        high-wmark 300 low-wmark 199
    exit
```

The system allows the various AA-admit-deny statistics to be exported via XML according to the configured accounting policy on the system. SAM-A can then use these statistics to generate the right reports / alerts.

As a prerequisite, an accounting policy is configured for aa-admit-deny statistics:

```
configure log accounting-policy 5 record aa-admit-deny
```

Then, the operator can configure AA to export the statistics related to various firewall functions configured in the system, such as:.

```
configure application-assurance group 1:1 statistics aa-admit-deny
                  accounting-policy 5
                  collect-stats
                  session-filter-stats
                  policer-stats-resources
                  tcp-validate-stats
              exit
```

GTP and STCP admit deny stats are related to firewall deployment within a SeGW, which is not covered within the scope of this chapter.

Show Routine — AQP:

```
*A:PE-1# show application-assurance group 1:1 policy app-qos-policy 110

===============================================================================
Application QOS Policy Entry 110 (Default Subscriber Policy)
===============================================================================
Description : FW for managed opted-in subs
Admin State : in-service
Hits:       : 0 flows
Conflicts   : 0 flows

Match :
    Traffic Direction        : network-to-subscriber
    ASO Characteristics      :
        FW-Protection                  : eq ON
Action :
    Session Filter               : denyUnsolictedwMgntCntrl
===============================================================
```

Show Routines — Session Filter:

```
*A:PE-1# show application-assurance group 1:1 session-
filter                                                "denyUnsolictedwMgntCntrl"
===============================================================================
AA Session Filter Instance "denyUnsolictedwMgntCntrl"
===============================================================================
Description    : (Not Specified)
```

```
Default Action : deny
    Event Log  : (Not Specified)
AQP Entries:    :
          110
-------------------------------------------------------

-------------------------------------------------------------------------------
Filter Match Criteria
------------------------------------------------------- ------------------------
-------------------------------------------------------
Entry          : 10
Description    : allow ICMP access from ISP LAN1
IP Protocol    : icmp
Source IP      : 10.10.8.0/24
Action         : permit
    Event Log  : (Not Specified)
Hits:          : 0 flows
------------------------------------------------------- ------------------------
-------------------------------------------------------
Entry          : 20
Description    : allow ICMP access from ISP LAN2
IP Protocol    : icmp
Source IP      : 192.168.0.0/24
Action         : permit
    Event Log  : (Not Specified)
Hits:          : 0 flows
------------------------------------------------------- ------------------------
-------------------------------------------------------
Entry          : 30
Description    : allow all TCP (e.g. FTP/telnet)access from ISP LAN2
IP Protocol    : tcp
Source IP      : 192.168.0.113/320/24
Action         : permit
    Event Log  : (Not Specified)
Hits:          : 0 flows
-------------------------------------------------------
-------------------------------------------------------------------------------
Entry          : 40
Description    : allow TCP on port 80 /HTTP access from ISP LAN1
IP Protocol    : tcp
Source IP      : 10.10.8.0/24
SourceDest Port      : eq 80
Action         : permit
    Event Log  : (Not Specified)
Hits:          : 0 flows
------------------------------------------------------- ------------------------
-------------------------------------------------------
No. of entries   : 4
=======================================================
```

Show Routines — TCP Validation:

```
*A:PE-1# show application-ass group 1:1 tcp-validate "TCP_protect"
===============================================================================
Application Assurance Group 1:1 tcp-validate "TCP_protect"
===============================================================================
Description      : (Not Specified)
Event log        : (Not Specified)
```

```
Strict Validation: No
AQP referenced   : Yes

-------------------------------------------------------------------------------
Decision Statistics                          sub-to-net              net-to-sub
-------------------------------------------------------------------------------
Total
-------------------------------------------------------------------------------
Allowed
  Octets                                              0                       0
  Packets                                             0                       0
Dropped
  Octets                                              0                       0
  Packets                                             0                       0

Dropped Reason
-------------------------------------------------------------------------------
Bad Flags
  Octets                                              0                       0
  Packets                                             0                       0
Bad Options
  Octets                                              0                       0
  Packets                                             0                       0
Bad Sequence Number
  Octets                                              0                       0
  Packets                                             0                       0
Bad Acknowledgment Number
  Octets                                              0                       0
  Packets                                             0                       0
No Establishment
    Octets                                            0                       0
  Packets                                             0                       0
SYN After Conn Establishment
  Octets                                              0                       0
  Packets                                             0                       0
Asymmetric Traffic
  Octets                                              0                       0
  Packets                                             0                       0
Traffic After Reset (RST)
  Octets                                              0                       0
  Packets                                             0                       0
Fragmented
  Octets                                              0                       0
  Packets                                             0                       0


*A:PE-1# show application-assurance threshold-crossing-alert detail
===============================================================================
Application Assurance Threshold Crossing Alerts
===============================================================================
-------------------------------------------------------------------------------
policer "Dos_police_Flow_count" from-sub
-------------------------------------------------------------------------------
Group:Part     : 1:1                    Trigger on     : denied-traffic
High watermark : 300                    Low watermark  : 199
Last raised    : N/A                    Last cleared   : N/A
State          : cleared
-------------------------------------------------------------------------------
tcp-validate "TCP_protect" from-sub
```

```
-------------------------------------------------------------------------------
Group:Part     : 1:1                    Trigger on     : denied-traffic
High watermark : 50                     Low watermark  : 40
Last raised    : N/A                    Last cleared   : N/A
State          : cleared
No. of TCAs : 2
===============================================================================
*A:PE-1#

*A:PE-1>tools>dump>app-assure>group# admit-deny-stats
===============================================================================
Application-Assurance Group 1:1 Admit-Deny Statistics
===============================================================================
-------------------------------------------------------------------------------
Admitted Sub-To-Net    Denied Sub-To-Net  Admitted Net-To-Sub    Denied Net-To-Sub
Packet Validation Statistics
         (Packets)            (Packets)            (Packets)            (Packets)
-------------------------------------------------------------------------------
Error
               0                    0                    0                    0
Fragments: Out-Of-Order
               0                    0                    0                    0
Fragments: All
               0                    0                    0                    0
Overload
             N/A                    0                  N/A                    0
-------------------------------------------------------------------------------

Admitted Sub-To-Net    Denied Sub-To-Net  Admitted Net-To-Sub    Denied Net-To-Sub
Session Filter Statistics
        (Sessions)           (Packets)           (Sessions)           (Packets)
-------------------------------------------------------------------------------
Session Filter: test
 Entry: 1                                                                    0
             0                    0                    0
 Default Action                                                              0
             0                    0                    0
-------------------------------------------------------------------------------

Admitted Sub-To-Net    Denied Sub-To-Net  Admitted Net-To-Sub    Denied Net-To-Sub
TCP Validation Statistics
         (Packets)            (Packets)            (Packets)            (Packets)
-------------------------------------------------------------------------------
test                                                                         0
             0                    0                    0
TCP_protect                                                                  0
             0                    0                    0
TCP_protect_ISP1                                                             0
             0                    0                    0
-------------------------------------------------------------------------------
*A:PE-1>tools>dump>app-assure>group#
```

Conclusion

The AA stateful packet filtering feature combined with AA Layer 7 classification and control empowers operators with an advanced, next generation firewall functionality that is integrated within SR OS. This chapter focused on traditional stateful and stateless session firewall functionality.

Application Assurance — Usage Monitoring and Policy Control via Diameter Gx Protocol

This chapter provides information about the diameter (Gx) control feature.

Topics in this chapter include:

Applicability

This configuration note is applicable to all 7750 SR/SR-c and 7450 ESS chassis supporting Application Assurance (AA).

The configuration was tested on release 13.0.R1.

Overview

The Gx reference point is defined in the Policy and Charging Control (PCC) architecture within the 3rd Generation Partnership Project (3GPP) standardization body. The Gx reference point is used for policy and charging control. The PCC architecture is defined in the 23.203 3GPP technical specification, while the Gx functionality is defined in the 29.212 3GPP technical specification. The SR OS implementation of Gx supports both Release 11 and Release 12 of the specification. Gx is an application of the Diameter protocol (RFC 6733). The Diameter protocol in SR OS is based on RFC 3588, *Diameter Base* Protocol.

As shown in Figure 26, Gx is placed between a policy server Policy and Charging Rule Function (PCRF) and a traffic forwarding node Policy and Charging Enforcement Function (PCEF) that enforces rules set by the policy server.

Figure 26 **Gx Reference Point**

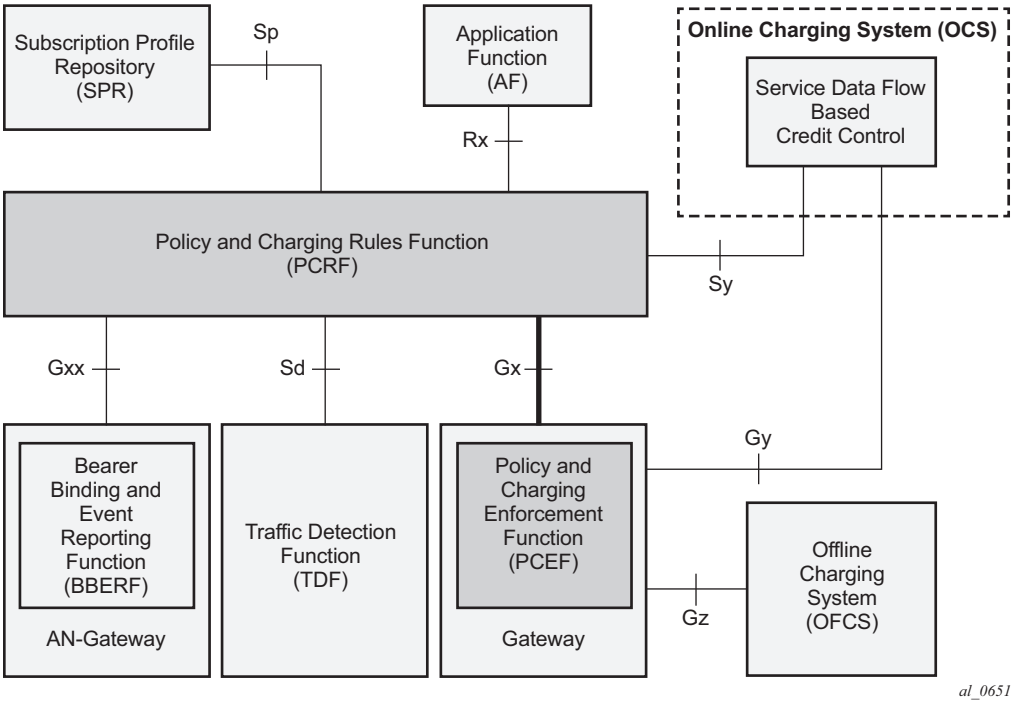

Although the Gx reference point is defined within the 3GPP standardization body, its applicability has also spread to wire-line operations to achieve mobile–fixed convergence gains by streamlining policy management functions into a single Gx based infrastructure, see Figure 27.

Figure 27 **Convergence**

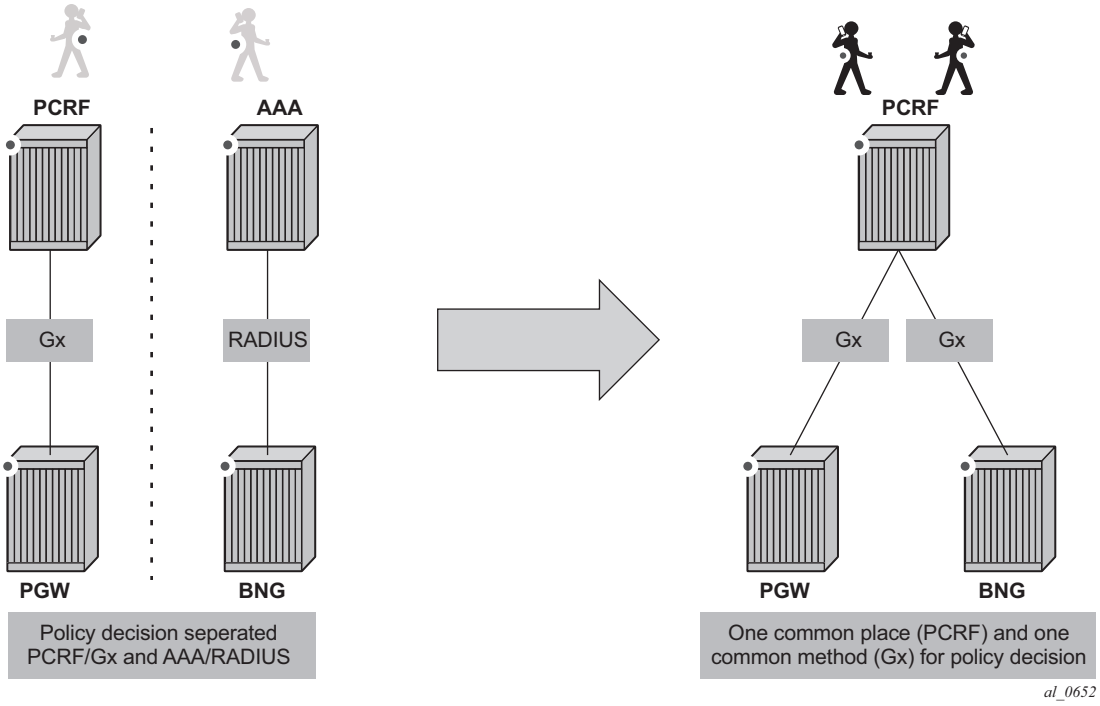

Gx support on SR OS is applicable to Enhanced Subscriber Management (ESM) functions, including the Application Assurance (AA) functions. The focus of this chapter is on the AA aspects of Gx.

The SR OS based Gx interface offers the following functionalities:

- ESM subscriber based policy decision providing
  - QoS attributes
  - charging attributes
  - subscriber identification
- Usage management
  - usage reporting from PCEF to PCRF

Figure 28 **Gx Reference Point**

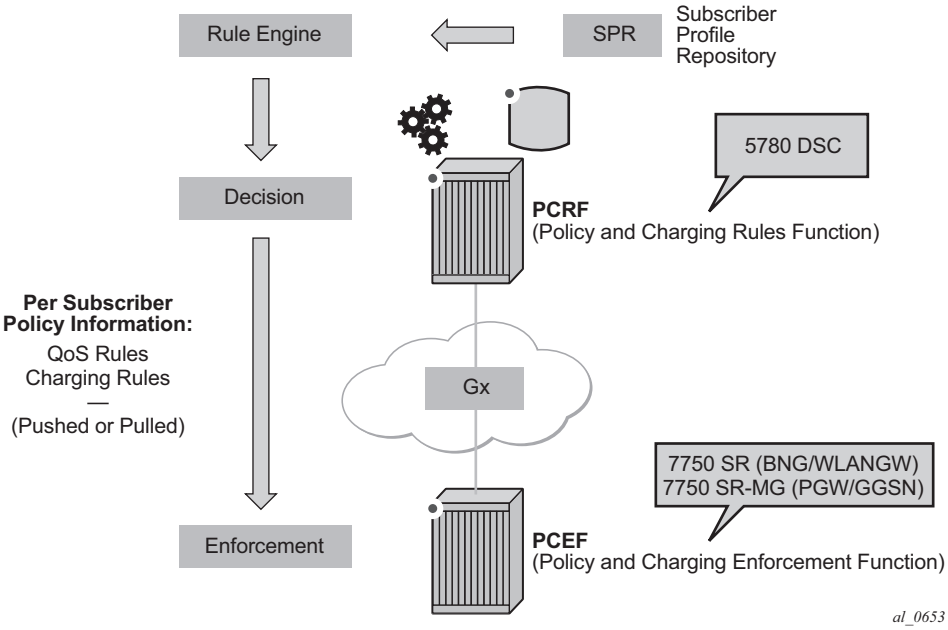

Note that Gx does not provide subscriber authentication or subscriber IP address assignment.

Figure 29 **Diameter Protocol Stack**

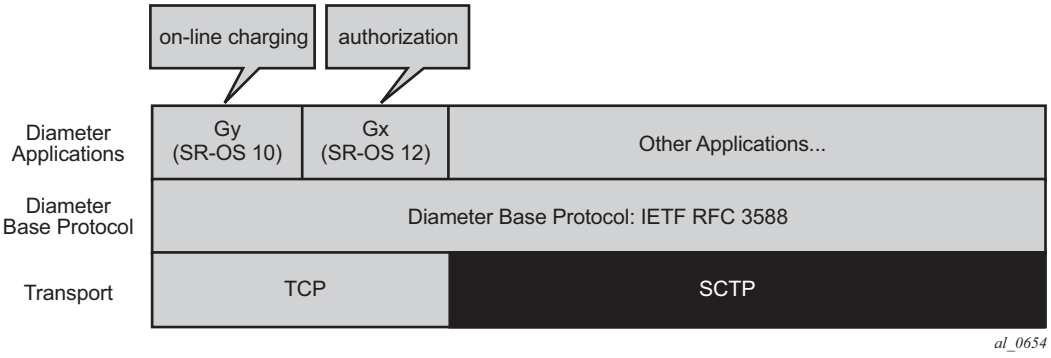

Policy Assignment Use Case

The SR OS accepts the following policy information from PCRF using Gx:

- Subscriber Profile strings and SLA Profile strings.

- Subscriber-QoS-Overrides.
- Application Profile strings.
- Application Subscriber Options (ASOs) related to AA.

Gx operates at subscriber host level and creates an “IP-CAN Session” (IP Connectivity Access Network) for every subscriber host. However, as AA operates at the subscriber level, AA related policies apply to all of the hosts belonging to that subscriber.

This chapter covers AA related functionalities, namely: application profile and ASO assignments and override. These functionalities are defined in either:

a. Application Detection and Control (ADC) rules -per 3GPP release 11 **or**
b. Policy and Charging Control (PCC) rules -per 3GPP release 12**- .**
   - **Application Profile** Alc-AA-Profile-Name Attribute-Value-Pair (AVP)
     - RADIUS equivalent is Alc-App-Prof-Str Vendor-Specific-Attribute (VSA)
   - **ASO overrides** Alc-AA-Service-Options AVP
     - RADIUS equivalent is Alc-AA-App-Service-Options VSA

Details of the ADC rules and related Nokia defined AVPs defined for use by AA are shown in Figure 30.

Figure 30 **ADC Rules and Related Nokia Defined AVPs Defined for Use by AA**

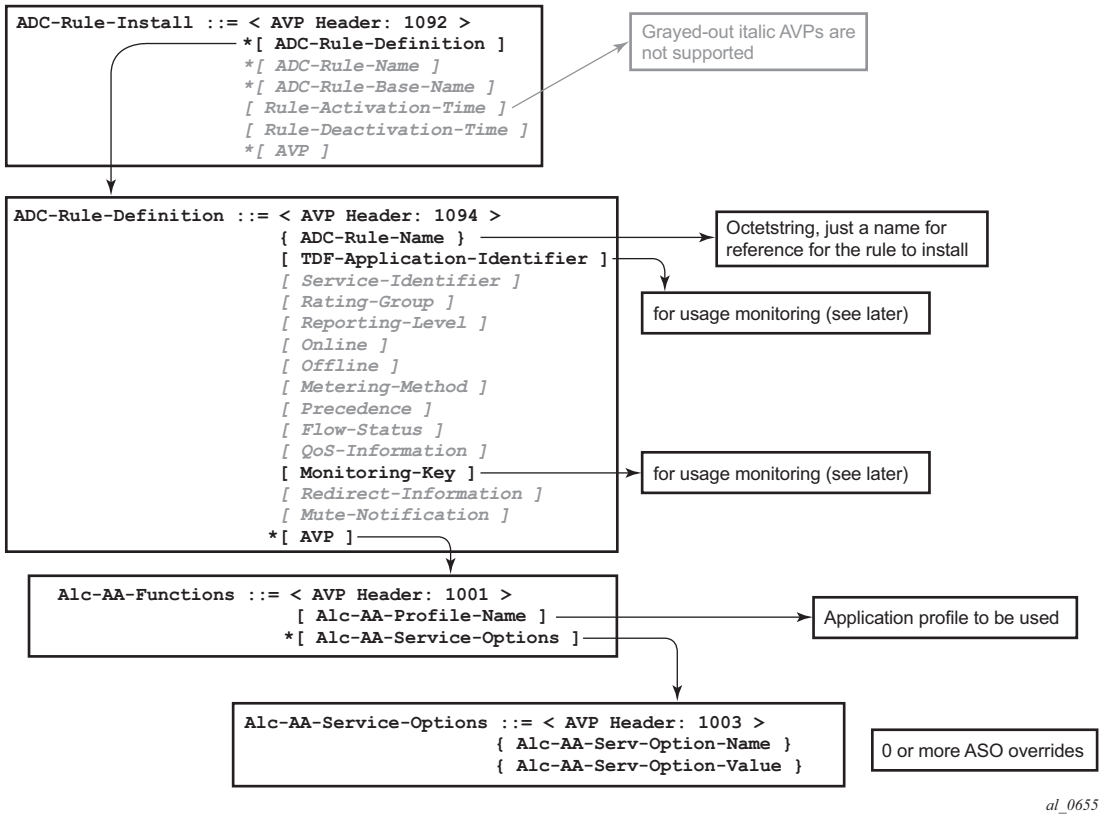

The ADC-Rule-Install is at the root level of the GX message.

As for 3GPP release 12, the details of the PCC rules and related Nokia defined AVPs defined for use by AA are shown in Figure 30.

Figure 31 **PCC Rules and Related Nokia-Defined AVPs Defined for Use by AA**

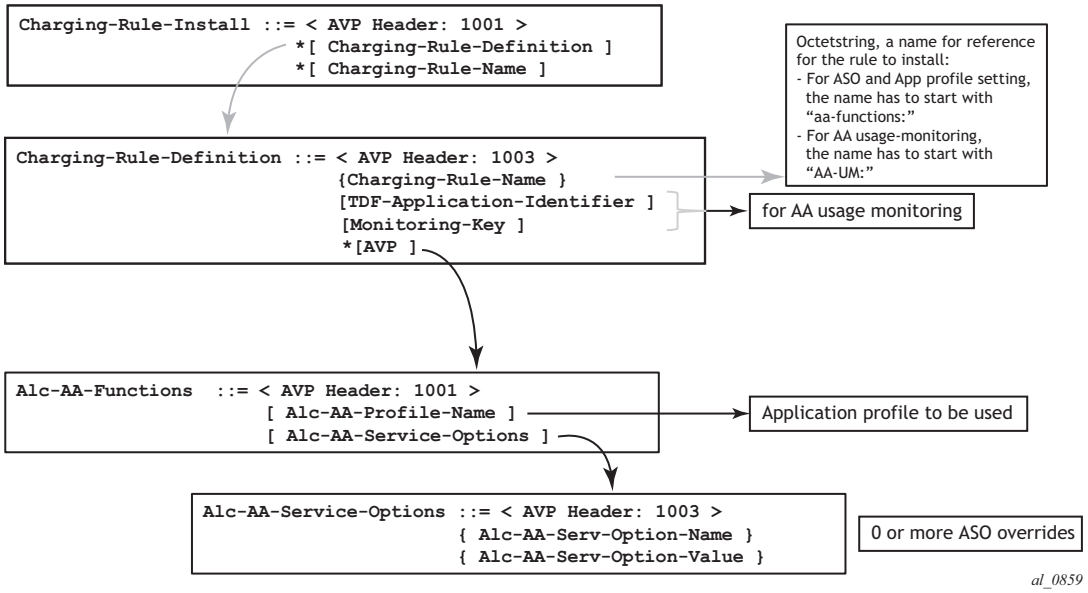

The PCC-Rule-Install, as in the case of ADC-Rule-Install, is at the root level of the GX message.

An example of the AVPs to install the application profile "gold_level" using 3GPP release 11 (/ADC rules) is shown in Figure 31.

Figure 32 **ADC Rule Example of AVPs to Install the Application Profile "gold_level"**

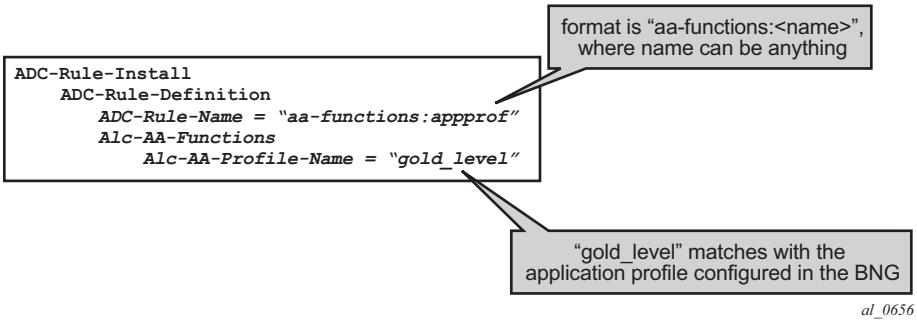

An example of the AVPs to install the "gold_level" application profile using 3GPP release 12 (/PCC rules) is shown in Figure 34.

Figure 33 **PCC Rule Example of AVPs to Install the Application Profile "gold_level"**

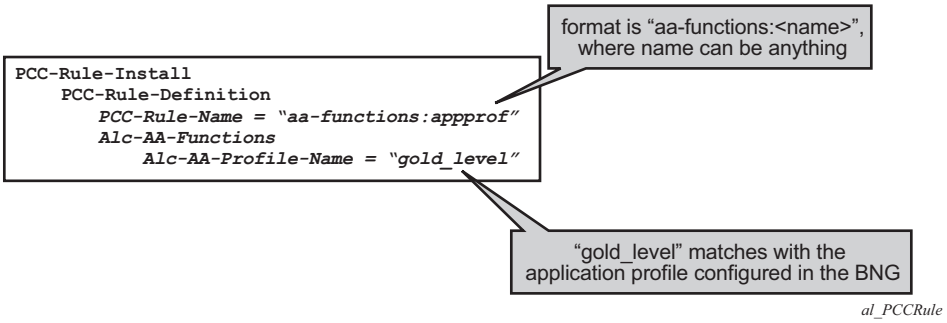

Note: ADC-Rule-Names and PCC-Rule-Names have to start with **aa-functions** when they contain an Alc-AA-Functions AVP.

The assignment of the **gold_level** appProfile is shown in another format in Figure 34.

Figure 34 **Capture of the ADC Rule Assignment of the "gold_level" appProfile**

```
adc-rule-install (1092) V------- [184]
       vendor-id TGPP
       data [172] (Grouped)
          adc-rule-definition (1094) V------ [172]
             vendor-id TGPP
             data [160] (Grouped)
               adc-rule-name (1096) V------ [32]
                 vendore-id TGPP
                 data [20] (UTF8String) : aa-functions:appprof
               AA-Functions (1001) V------ [128]
                 vendor-id ALU
                 data [116] (Grouped)
                   AA-Profile-Name (1002) V------ [17]
                   vendor-id ALU
                   data [5] (UTF8String) : gold level
                   AA-App-Service-Options (1003) V------ [48]
                     vendor-id ALU
                     data [36] (Grouped)
                       AA-App-Service-Options-Name (1004) V------ [17]
                         vendor-id ALU
                         data [5] (UTF8String) : level
                       AA-App-Serv-Options-Value (1005) V------ [16]
                         vendor-id ALU
                         data [4] (UTF8String) : high
                   AA-App-Service-Options (1003) V------ [48]
                     vendor-id ALU
                     data [36] (Grouped)
                       AA-App-Serv-Options-Name (1004) V------ [18]
                         vendor-id ALU
                         data [6] (UTF8String) : p2p
                       AA-App-Serv-Options-Value (1005) V------ [14]
                         vendor-id ALU
                         data [2] (UTF8String) : unlimited
```

al_0657

Application profiles and ASO overrides can be changed on-the-fly with a Re-Authentication-Request (RAR) message according to these rules:

- If an Application profile is present in the Gx message it is applied first. Then ASO AVPs are applied when present (in the Gx message). In other words:
  - If a RAR message only contains the same application profile and no ASO overrides, then all previous ASO overrides are removed.
  - When a RAR message contains the same application profile and new ASO overrides, then the new ASO overrides are applied, and the previous ASO overrides are removed.
  - When a RAR message contains a new application profile, all previous ASO overrides are removed and replaced with the ASOs in the RAR if present.
  - When a RAR message does not contain an application profile but only ASO overrides, then the new ASO overrides are added to the existing ASO overrides.

Note that a single Gx ADC (or PCC) rule cannot contain both AA subscriber policies (appProfile/ASO) and AA Usage monitoring (as outlined later). These have to be in separate ADC (or PCC) rules.

Usage Management/Monitoring Use-Case

The AA-ISA can monitor application usage at the subscriber level and report back to the PCRF whenever the usage exceeds the threshold(s) set by the PCRF when receiving requests from the PCRF over the Gx interface.

Usage monitoring can be used by operators to report to PCRF when:

- The AA-ISA detects the start of a subscriber application by setting the usage threshold to a very low value.
- A pre-set usage volume per subscriber application is exceeded.

AA can monitor subscriber's traffic for any defined:

- Application,
- Application group, and/or
- Charging group.

The AA-ISA reports the accumulated usage when:

- A usage threshold is reached.
- The PCRF explicitly disables the usage monitoring.
- The PCRF requests a report.
- The ADC (or PCC) rule associated with the monitoring instance is removed or deactivated.
- A session is terminated.

An AA defined application, application group and/or charging group is automatically allowed to be referenced by an ADC (or PCC) rule for the purpose of usage monitoring only if:

{It is already selected for either XML or RADIUS per subscriber accounting

OR

It is explicitly enabled by the operator for per subscriber statistics collection}

AND

Usage monitoring is enabled for the given AA group:partition

Figure 35 illustrates the different messaging/call flows involved in application level usage monitoring. Details of the different supported AVPs used in these messages are illustrated later.

Figure 35 **Call Flow Diagram**

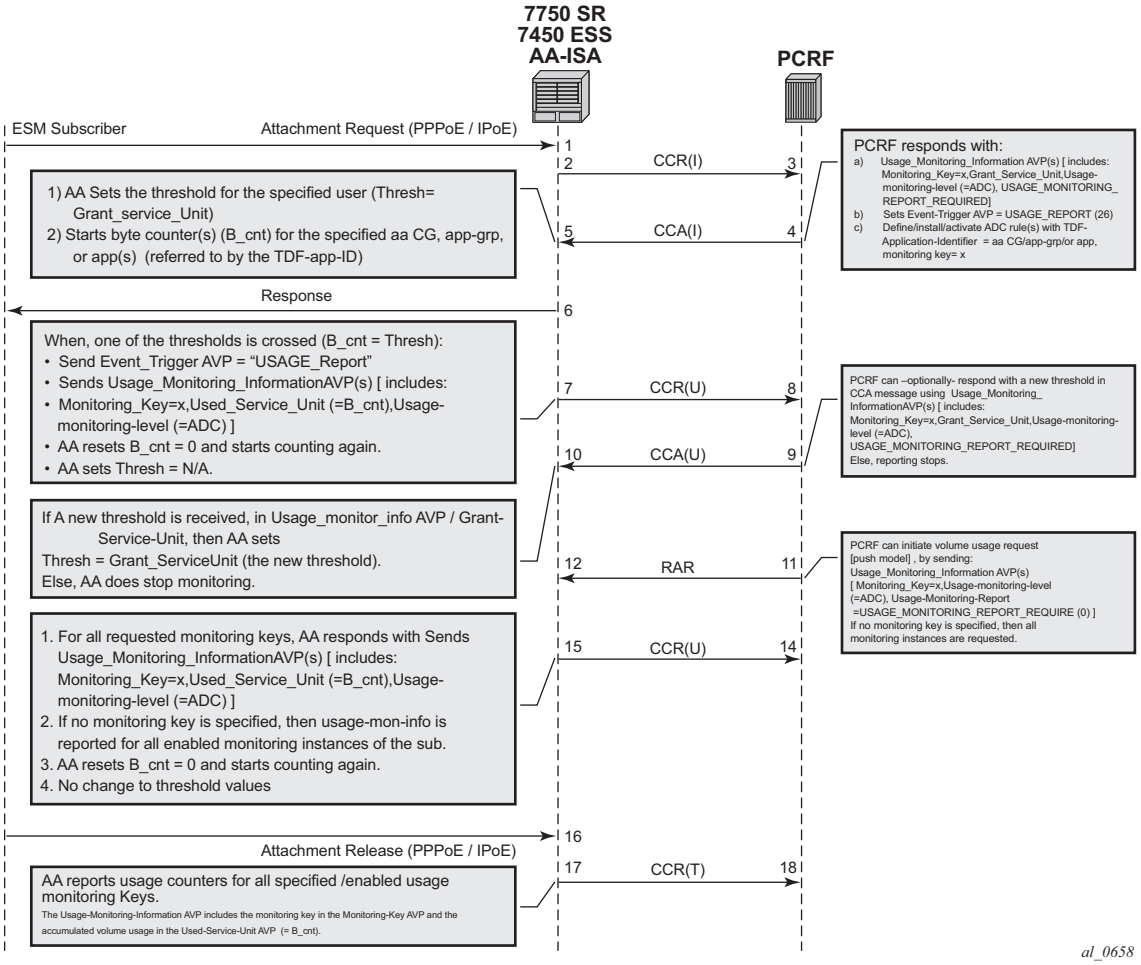

The AA-ISA/PCEF supports Usage-Thresholds AVPs that refer to the thresholds (in bytes) at which point an event needs to be sent back to the PCRF, (see Figure 33).

Time based thresholds are not supported.

AA supports the "grant-service-unit" AVP using the following possible values (AVP):

- CC-Input-Octets AVP (code 412) : From Subscriber total byte count threshold.
- CC-Output-Octet AVP (code 414): To subscriber total byte count threshold.
- CC-Total-octets AVP (code 421): Threshold of aggregate traffic (Input and Output byte counters).

As shown in Figure 35, (T=7), AA sends a Credit Control Request (CCR_ message) with a "USAGE_REPORT" Event-Trigger AVP to the PCRF when the usage counter reaches the configured usage monitoring threshold for a given subscriber (and given application group). AA counters are reset (to zero) when the monitoring threshold is reached (and an event is sent back to the PCRF). The counter(s) however does not stop counting newly arriving traffic. AA counters only include “admitted” packets. Any packets that were discarded by AA due to, for example, policing actions are not counted for usage-monitoring purposes.

The TDF-Application-Identifier AVP (within the ADC or PCC rule) refers to an AA Charging group, an AA application group or to an AA application. TDF-Application-Identifiers (for example, charging-groups) have to be manually entered at the PCRF to match the AA charging groups defined in the AA. If the TDF-Application-Identifier refers to a name that is used for both a charging group and an application (or an application group), AA monitors the charging group. In other words, the AA charging group has a higher precedence than the AA application group.

Gx Usage Monitoring AVP Summary

For 3GPP release 11 (using ADC rules), the following AVPs are used for AA-Usage monitoring:

```
ADC-Rule-Install ::= < AVP Header: 1092 >
                            *[ ADC-Rule-Definition ]
                            *[ ADC-Rule-Name ]


    ADC-Rule-Definition ::= < AVP Header: 1094 >
                                  { ADC-Rule-Name }
                                  [ TDF-Application-Identifier ]; AA app/app-grp/
chrg-grp
                                  [ Monitoring-Key  ];


    Usage-Monitoring-Information::= < AVP Header: 1067 >
                                    [ Monitoring-Key ]
                                 0,2[ Granted-Service-Unit ]
                                        Granted-Service-Unit ::= < AVP Header: 431 >
                                    [ CC-Total-Octets ]
                                    [ CC-Input-Octets ]
                                    [ CC-Output-Octets ]

                               0,2[ Used-Service-Unit ]
    Used-Service-Unit ::= < AVP Header: 446 >
                            [ CC-Total-Octets ] ;
                            [ CC-Input-Octets ]
                            [ CC-Output-Octets ]

                                 [ Usage-Monitoring-Level ]
; ADC_RULE_LEVEL (2)
```

```
                                [ Usage-Monitoring-Report ]
; immidiate report -- USAGE_MONITORING_REPORT_REQUIRED (0)

                                [ Usage-Monitoring-Support ]
; to disable :  USAGE_MONITORING_DISABLED (0)
```

For 3GPP release 12 (using PCC rules), the following AVPs are used for AA-Usage monitoring:

```
Charging-Rule-Install ::= < AVP Header: 1001 >
                             *[ Charging-Rule-Definition ]
                             *[ Charging-Rule-Name ]


    Charging-Rule-Definition ::= < AVP Header: 1003 >
                                  { Charging-Rule-Name } ;/ starts with "UM-AA:"
                        [ TDF-Application-Identifier ]; AA app/app-grp/chrg-grp
                        [ Monitoring-Key ];



    Usage-Monitoring-Information::= < AVP Header: 1067 >
                           [ Monitoring-Key ]
                         0,2[ Granted-Service-Unit ]
                               Granted-Service-Unit ::= < AVP Header: 431 >
                                [ CC-Total-Octets ]
                                [ CC-Input-Octets ]
                                [ CC-Output-Octets ]

                         0,2[ Used-Service-Unit ]
    Used-Service-Unit ::= < AVP Header: 446 >
                           [ CC-Total-Octets ] ;
                           [ CC-Input-Octets ]
                           [ CC-Output-Octets ]

                            [ Usage-Monitoring-Level ]
; PCC_RULE_LEVEL (1)

                           [ Usage-Monitoring-Report ]
; immidiate report -- USAGE_MONITORING_REPORT_REQUIRED (0)

                           [ Usage-Monitoring-Support ]
; to disable :  USAGE_MONITORING_DISABLED (0)
```

Configuration

This configuration example highlights the commands illustrating how Gx can be used to:

- Override AppProfile and ASO characteristics.
- Set and retrieve AA level usage monitoring metrics.

While the configuration associated with setting up the Gx interface toward the PCRF is shown for the sake of completeness, that aspect of the configuration is not explored in detail, and only a 3GPP release 11 model is shown Similarly, the Gx policies and usage monitoring associated with ESM host policies (non-AA aspects) are out of the scope of this chapter.

The configuration on the 7750 node is the same, independent of whether the PCRF supports 3GPP release 11 (ADC) or release 12(PCC) to provide AA policy control function.

Figure 36 **Example Configuration Setup**

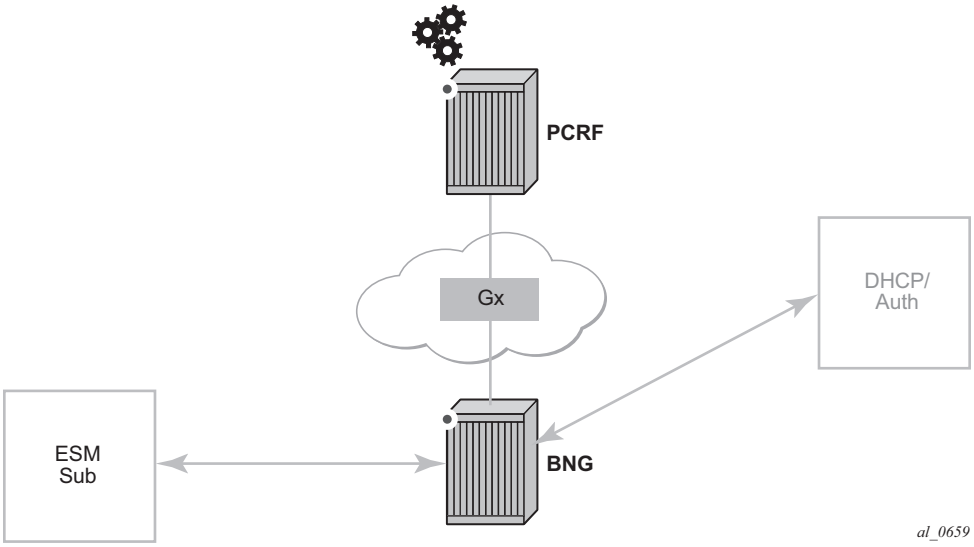

The BNG is set up with at least one IOM and one MS-ISA MDA configured as ISA-AA.

```
configure
    card 1
        card-type iom3-xp
        mda 1
            mda-type m20-1gb-xp-sfp
            no shutdown
```

```
        exit
        mda 2
            mda-type isa-aa
            no shutdown
        exit
        no shutdown
    exit
    card 3
        card-type iom3-xp
        mda 1
            mda-type isa-aa
            no shutdown
        exit
        mda 2
            mda-type isa-aa
            no shutdown
        exit
        no shutdown
    exit
```

The configurations in this example are broken down into four main steps:

Step 1. Configuring the Gx interface (high-level)

Step 2. Configuring AA application profiles and ASOs (high-level)

Step 3. Configuring AA applications filters (high-level)

Step 4. Configuring AA usage-monitoring

The focus of this configuration example is on Step 4, and the updated show routines related to AA ESM subscriber state are shown at the end of Step 2.

Step 1. Configuring the Gx interface (high-level).

These commands bring up the Gx diameter control channel between the Gx Controller(/Server), also known as PCRF, and the PCEF(/BNG).

```
configure
 aaa
    diameter-peer-policy "ppol" create
       applications gx
        connection-timer 5
        origin-host "router.workstation"
        origin-realm "Nokia.com"
        transaction-timer 5
        watchdog-timer 10
        peer "ppeer0" create
            address 10.1.0.10
            destination-host "primary-pcrf.Nokia.com"
            destination-realm "Nokia.com"
            no shutdown
        exit
    exit
exit
```

The diameter peer policy “**ppol**” is then referenced under subscriber management.

```
configure
  subscriber-mgmt
    diameter-application-policy "diamAppPlcy" create
        application gx
        diameter-peer-policy "ppol"
    exit
```

Then the created subscriber management policy “**diamAppPlcy**” is applied to the subscriber interface.

```
configure
    service
        customer 1 create
            description "Default customer"
        exit
        ies 1 customer 1 vpn 1 create
            description "Default Ies description for service id 1"
            subscriber-interface "ies-1-172.16.0.0" create
                address 172.16.0.0/12
                group-interface "grp-1-35782656-1" create
                    dhcp
                       server 172.16.200.200
                       trusted
                       lease-populate 2000
                       gi-address 172.16.0.0
                       no shutdown
                    exit
                    diameter-application-policy "diamAppPlcy"
                    sap 1/1/4:1 create
                     description "sap-grp-1"
                         sub-sla-mgmt
                             def-sub-profile "sub_prof"
                             def-sla-profile "sla_prof"
                             def-app-profile "app_prof_1"
                             sub-ident-policy "sub_ident_A_1"
                             multi-sub-sap 2
                             no shutdown
                         exit
                     exit
                exit
            exit
            service-name "ACG Ies 1"
            no shutdown
        exit
```

Now verify the configuration and connectivity towards the PCRF by running the following command:

```
*A:BNG-1# show aaa diameter-peer-policy "ppol"
===============================================================================
Diameter Peer Policy : ppol
===============================================================================
Last Mgmt Change      : 05/30/2014 18:53:38
```

```
-------------------------------------------------------------------------------
Diameter Config Values
-------------------------------------------------------------------------------
Origin Host           : router.workstation.be
Origin Realm          : lucent.com
Connection Timer      : 5                  Source Address     : 0.0.0.0
Transaction Timer     : 5                  Router             : Base
Watchdog Timer        : 10
Vendor Support        : 3GPP (default)
Python Policy         : N/A
-------------------------------------------------------------------------------
Peer Name            Oper  PSM State     Susp  Cooldown  Pref  Order  Pri/Sec
-------------------------------------------------------------------------------
ppeer0               Yes   I-Open        No    -         50    1      Primary
===============================================================================
*A:BNG-1#
```

The Peer-State-Machine State (PSM), as per RFC 3588, has the value I-OPEN indicating that the peer is operational. The “I-” stands for **I**nitiator state, in this case the BNG is the initiator.

A detailed look into the traffic statistics between the PCEF and the PCRF (Gx controller) can be viewed using a show statistics command (see below). These statistics provide a breakdown of the messages exchanged:

```
*A:BNG-1# show aaa diameter-peer-policy "ppol" peer "ppeer0" statistics
===============================================================================
Diameter Peer Policy : ppol (statistics)
===============================================================================
Diameter Peer             : ppeer0
time statistics cleared   : 05/30/2014 18:53:38
-------------------------------------------------------------------------------
Client initiated tx/rx                   Server initiated tx/rx
-------------------------------------------------------------------------------
TCP Send Failed           : 0            TCP Send Failed           : 0
Diam Rx Drop Count (Resps): 0            Diam Rx Drop Count (Reqs) : 0
Diam Tx Requests          : 313          Diam Rx Requests          : 204
Diam Rx Responses         : 313          Diam Tx Responses         : 204
Pending Messages          : 0
Request Timeouts          : 0
-------------------------------------------------------------------------------
Diameter message breakdown
-------------------------------------------------------------------------------
CCR initial Tx            : 111          CCA initial Rx            : 111
CCR update Tx             : 88           CCA update Rx             : 88
CCR terminate Tx          : 11           CCA terminate Rx          : 11
CER Tx                    : 1            CEA Rx                    : 1
DWR Tx                    : 102          DWA Rx                    : 102
DWR Rx                    : 0            DWA Tx                    : 0
ASR Rx                    : 0            ASA Tx                    : 0
RAR Rx                    : 204          RAA Tx                    : 204
DPR Tx                    : 0            DPA Rx                    : 0
DPR Rx                    : 0            DPA Tx                    : 0
===============================================================================
*A:BNG-1#
```

Step 2. Configuring AA application profiles and ASOs (high-level)

To illustrate the use of application profiles and ASO overrides using Gx RAR messages, four ASOs and 2 appProfiles are defined, see below.

“app_prof_1“ is the default app-profile used when a subscriber is created on AA.

```
configure
   application-assurance
       group 129:34883 create
           policy
               begin
               app-service-options
                   characteristic "permitDNS" persist-id 1 create
                       value "no"
                       value "yes"
                       default-value "yes"
                   exit
                   characteristic "permitRDP" persist-id 2 create
                       value "no"
                       value "yes"
                       default-value "yes"
                   exit
                   characteristic "permitHTTP" persist-id 3 create
                       value "no"
                       value "yes"
                       default-value "yes"
                   exit
               exit
               app-profile "app_prof_1" create
                   description "Application Profile Id app_prof_1"
                   divert
               exit
               app-profile "app_prof_2" create
                   description "Application Profile Id app_prof_2"
                   divert
               exit
```

Step 3. Configuring AA applications filters (high-level)

First create the application group, as follows.

```
configure isa
       application-assurance-group 129 create
           primary 3/2
           backup 1/2
           partitions
           divert-fc be
           no shutdown
       exit
```

Then create the partition and associated charging groups, application groups, applications, etc.

```
configure
    application-assurance
        group 129:34883 create
            policy
```

```
                    begin
                    charging-group "0_rated" create
                        export-id 1
                    exit
                    charging-group "default_charge_group" create
                        export-id 255
                    exit
                    default-charging-group "default_charge_group"
                    app-group "Other" create
                        export-id 8
                    exit
                    app-group "Peer to Peer" create
                        export-id 3
                    exit
                    app-group "Remote Connectivity" create
                        export-id 4
                    exit
                    app-group "Unknown"
                        charging-group "0_rated"
                        export-id 1
                    exit
                    app-group "Web" create
                        export-id 10
                    exit
                    application "DNS" create
                        description "default-description for application DNS"
                        app-group "Other"
                        export-id 12
                    exit
                    application "BitTorrent" create
                        app-group "Peer to Peer"
                        export-id 3
                    exit
                    application "HTTP" create
                        description "default-description for application HTTP"
                        app-group "Web"
                        export-id 26
                    exit
                    application "RDP" create
                        description "default-description for application RDP"
                        app-group "Remote Connectivity"
                        export-id 61
                    exit
                    application "Unknown"
                        charging-group "0_rated"
                        export-id 1
                    exit
                exit
                commit
            exit
        exit
    exit
```

Example app-filter definitions defining HTTP, DNS, Bittorrent and RDP applications are show below.

```
configure
    application-assurance
        group 129:34883
```

```
            policy
                begin
                app-filter
                    entry 6 create
                        description "default-description for AppFilter entry 6"
                        protocol eq "rdp"
                        ip-protocol-num eq tcp
                        application "RDP"
                        no shutdown
                    exit
                    entry 9 create
                        description "default-description for AppFilter entry 9"
                        protocol eq "dns"
                        ip-protocol-num eq udp
                        server-port eq range 53 55
                        application "DNS"
                        no shutdown
                    exit
                    entry 20 create
                        description "default-description for AppFilter entry 20"
                        protocol eq "bittorrent"
                        ip-protocol-num eq tcp
                        application "BitTorrent"
                        no shutdown
                    exit
                    entry 38 create
                        description "default-description for AppFilter entry 38"
                        protocol eq "http"
                        ip-protocol-num eq tcp
                        server-port gt 8738
                        application "HTTP"
                        no shutdown
                    exit
                exit
                commit
            exit
        exit
    exit
```

Note: The focus of this example is on the definition of app-filters and/or AQPs. These are listed above (and below) for illustration purposes. The “sample” AQP configurations and app-filters shown here should not be used in a real-life configuration. Their configuration should follow the information in Application Assurance — Application Identification and User-Defined Applications.

Example AQP configurations for blocking DNS, RDP and HTTP traffic are listed below.

```
configure
    application-assurance
        group 129:34883
            policy
                begin
                app-qos-policy
                    entry 2 create
```

```
                        match
                            application eq "DNS"
                            characteristic "permitDNS" eq "no"
                        exit
                        action
                            drop
                        exit
                        no shutdown
                    exit
                    entry 3 create
                        match
                            application eq "HTTP"
                            characteristic "permitHTTP" eq "no"
                            ip-protocol-num neq 0
                        exit
                        action
                            drop
                        exit
                        no shutdown
                    exit
                    entry 4 create
                        match
                            application eq "RDP"
                            app-group eq "Remote Connectivity"
                            characteristic "permitRDP" eq "no"
                            ip-protocol-num neq udp
                        exit
                        action
                            drop
                        exit
                        no shutdown
                    exit
                exit
                commit
            exit
        exit
    exit
```

When an ESM subscriber is created, it is associated with the default AA app-profile, as seen using the show command below.

```
*A:BNG-1>show>app-assure>group# aa-sub esm "sub_172.16.0.2" summary
===============================================================================
Application-Assurance Subscriber Summary (realtime)
===============================================================================
AA-Subscriber           : sub_172.16.0.2 (esm)
ISA assigned            : 3/2
App-Profile             : app_prof_1
App-Profile divert      : Yes
Capacity cost           : 1
Aarp Instance Id        : N/A
HTTP URL Parameters     : (Not Specified)
Last HTTP Notified Time : N/A
-------------------------------------------------------------------------------
Traffic                         Octets              Packets                Flows
-------------------------------------------------------------------------------
From subscriber:
  Admitted                           0                    0                    0
```

```
  Denied                              0                   0                   0
  Active flows                                                                0
To subscriber:
  Admitted                            0                   0                   0
  Denied                              0                   0                   0
  Active flows                                                                0
Flow counts:
  Terminated                                                                  0
  Short duration                                                              0
  Med duration                                                                0
  Long duration                                                               0
Total flow duration :  0 seconds
-------------------------------------------------------------------------------
Top App-Groups                                Octets        Packets        Flows
-------------------------------------------------------------------------------
None
-------------------------------------------------------------------------------
Application Service Options (ASO)
-------------------------------------------------------------------------------
Characteristic                  Value                              Derived from
-------------------------------------------------------------------------------
permitDNS                       yes                                default
permitRDP                       yes                                default
permitHTTP                      yes                                default
===============================================================================
*A:BNG-1>show>app-assure>group#
```

After the PCRF sends out AppProfile and ASO override AVPs, using either PCC or ADC rules, in RAR messages (as shown below) it can be seen that the new parameters (new profile and new values for permitDNS and permitHTTP ASOs) are updated for that ESM subscriber.

Figure 37 **PCRF AVPs Override Call Flow Diagram**

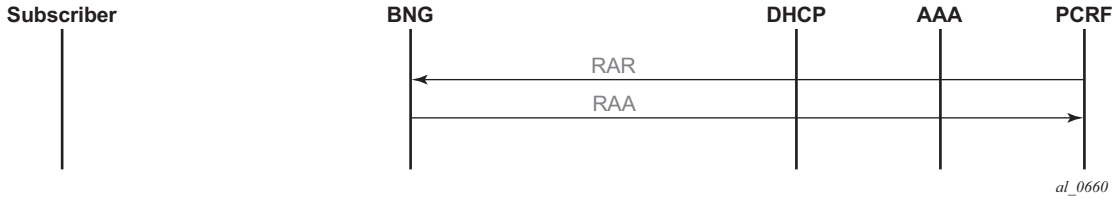

Figure 38 **RAR Containing ASOs and AppProfile Override AVPs Example**

```
adc-rule-install (1092) V------- [184]
       vendor-id TGPP
       data [172] (Grouped)
          adc-rule-definition (1094) V------ [172]
             vendor-id TGPP
             data [160] (Grouped)
               adc-rule-name (1096) V------ [32]
                 vendore-id TGPP
                 data [20] (UTF8String) : aa-functions:appprof
               AA-Functions (1001) V------ [128]
                 vendor-id ALU
                 data [116] (Grouped)
                   AA-Profile-Name (1002) V------ [17]
                   vendor-id ALU
                   data [5] (UTF8String) : app_prof 2
                   AA-App-Service-Options (1003) V------ [48]
                     vendor-id ALU
                     data [36] (Grouped)
                       AA-App-Serv-Options-Name (1004) V------ [17]
                         vendor-id ALU
                         data [5] (UTF8String) : permitDNS
                       AA-App-Serv-Options-Value (1005) V------ [16]
                         vendor-id ALU
                         data [4] (UTF8String) : no
                 AA-App-Service-Options (1003) V------ [48]
                   vendor-id ALU
                   data [36] (Grouped)
                     AA-App-Serv-Options-Name (1004) V------ [18]
                       vendor-id ALU
                       data [6] (UTF8String) : permitHTTP
                     AA-App-Serv-Options-Value (1005) V------ [14]
                       vendor-id ALU
                       data [2] (UTF8String) : no
```

al_0661

```
*A:BNG-1>show>app-assure>group# aa-sub esm "sub_172.16.0.2" summary
===============================================================================
Application-Assurance Subscriber Summary (realtime)
===============================================================================
AA-Subscriber           : sub_172.16.0.2 (esm)
ISA assigned            : 3/2
App-Profile             : app_prof_2
App-Profile divert      : Yes
Capacity cost           : 1
Aarp Instance Id        : N/A
HTTP URL Parameters     : (Not Specified)
Last HTTP Notified Time : N/A
-------------------------------------------------------------------------------
Traffic                          Octets              Packets                Flows
-------------------------------------------------------------------------------
From subscriber:
  Admitted                            0                    0                    0
  Denied                              0                    0                    0
  Active flows                                                                  0
To subscriber:
  Admitted                            0                    0                    0
  Denied                              0                    0                    0
  Active flows                                                                  0
Flow counts:
  Terminated                                                                    0
  Short duration                                                                0
  Med duration                                                                  0
  Long duration                                                                 0
```

```
Total flow duration :  0 seconds

-------------------------------------------------------------------------------
Top App-Groups                                Octets        Packets        Flows
-------------------------------------------------------------------------------
None

-------------------------------------------------------------------------------
Application Service Options (ASO)
-------------------------------------------------------------------------------
Characteristic                Value                              Derived from
-------------------------------------------------------------------------------
permitDNS                     no                                 dyn-override
permitRDP                     yes                                default
permitHTTP                    no                                 dyn-override
===============================================================================
```

Step 4. Configuring AA Usage Monitoring

Once the applications, application groups and/or charging groups are defined and configured (see previous steps), the operator needs:

- to enable the collection of per-subscriber statistics so they can be used for Gx based usage-monitoring. This step is not needed for any app/appgrp or charging group that is already enabled for per-subscriber statistics. In other words, if XML or RADIUS accounting is enabled for a given app/appgrp or charging group, then Gx usage-monitoring is also automatically enabled.
- to enable usage-monitoring for the given AA group:partition.

```
config
    application-assurance
        group 129:34883
            statistics
                aa-sub
                    usage-monitoring
                    app-group "Unknown" export-using accounting-policy
                                                    radius-accounting-policy
                    charging-group "0_rated" export-using accounting-policy
                                                    radius-accounting-policy
                    charging-group "default_charge_group" export-using
                                                          accounting-policy
                    radius-accounting-policy
                    application "BitTorrent" no-export
                exit
```

In the example above:

- The usage-monitoring command is used to enable Gx usage monitoring for the specified AA partition.
- The aa-group and charging-group commands specify which charging groups and AA groups are selected for export. In this case *0-rated*, *Unknown*, and *default-charging-group* are selected for RADIUS accounting and they automatically qualify for Gx-usage monitoring.

- The BitTorrent application however needs to be explicitly configured as “no-export” as it needs to be enabled for Gx-usage monitoring.

The operator can display the number of usage monitoring rules for a given subscriber. This is shown below after the ESM subscriber is created, but before any ADC rules are installed for usage-monitoring by PCRF, so AA reports that no rules apply (“0”).

```
*A:BNG-1>show>app-assure>group# aa-sub esm "alcatel_A_1" usage-monitor status
===============================================================================
Application-Assurance Subscriber "alcatel_A_1" (esm)
Usage Monitor Status
===============================================================================
Type             Name                              Oper Status
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No. of rules: 0
===============================================================================
*A:BNG-1>show>app-assure>group#
```

The PCRF then sends a RAR message with a usage monitoring ADC or PCC rule for the BitTorrent application to set the usage thresholds for BitTorrent for the ESM subscriber “alcatel_A_1” to (in bytes):

Input (from sub) 1378168

Output (to sub) 1381148

Total traffic (up and down) 18446744073709551614

Figure 39 **RAR Containing Usage Monitoring ADC Rules Example**

```
adc-rule-install (1092) V------- [96]
       vendor-id TGPP
       data [84] (Grouped)
          adc-rule-definition (1094) V------ [84]
             vendor-id TGPP
             data [72] (Grouped)
               adc-rule-name (1096) V------ [20]
                  vendore-id TGPP
                  data [8] (UTF8String) whatever
               tdf-application-id (1088) V------[22]
                  vendor-id ALU
                  data [10] (UTF8String) : BitTorrent
               monitoring-key (1066) V------ [25]
                  vendor-id TGPP
                  data [13] (UTF8String) : torrentmonkey
```

```
usage-monitoring-information (1067) V------- [80]
       vendor-id TGPP
       data [68] (Grouped)
          monitoring-key (1066) V------ [25]
             vendor-id TGPP
             data [13] (UTF8String) : torrentmonkey
          granted-service-units (431) ------- [24]
             data [16] (Grouped)
               cc-input-octets (412) ------ [16]
                  data [8] (Unsigned64) : 1378168
               cc-output-octets (414) ------ [16]
                  data [8] (Unsigned64) : 1378168
               cc-total-octets (421) ------ [16]
                  data [8] (Unsigned64) : 18446744073709551614
          monitoring-key (1068) V------ [16]
             vendor-id TGPP
             data [4] (Enumerated) : 2 : ADC RULE LEVEL
```

al_0662

This is then reflected on the AA-ISA:

```
*A:BNG-1>show>app-assure>group# aa-sub esm "alcatel_A_1" usage-monitor status
===============================================================================
Application-Assurance Subscriber "alcatel_A_1" (esm)
Usage Monitor Status
===============================================================================
Type             Name                              Oper Status
-------------------------------------------------------------------------------
application      BitTorrent                        active
-------------------------------------------------------------------------------
No. of rules: 1
===============================================================================
*A:BNG-1>show>app-assure>group#
```

Note the “active” oper status is set since there is at least one usage monitoring threshold associated with this application.

Given that there is no traffic flowing yet to or from the subscriber the counters currently are “0”:

```
*A:BNG-1>show>app-assure>group# aa-sub esm "alcatel_A_1" usage-monitor count
===============================================================================
Application-Assurance Subscriber "alcatel_A_1" (esm)
Usage Monitor Credit Statistics
===============================================================================
Application: "BitTorrent"
Direction     Status                    Granted                 Used     % Used
-------------------------------------------------------------------------------
to sub        valid                     1378168                    0         0%
from sub      valid                     1381148                    0         0%
both          valid        18446744073709551614                    0         0%
===============================================================================
*A:BNG-1>show>app-assure>group#
```

The status is set to “valid” since a threshold (or Grant) is received.

When, at a later stage, traffic starts flowing again usage-monitor subscriber statistics are updated as shown below.

```
*A:BNG-1>show>app-assure>group# aa-sub esm "alcatel_A_1" usage-monitor count
===============================================================================
Application-Assurance Subscriber "alcatel_A_1" (esm)
Usage Monitor Credit Statistics
===============================================================================
Application: "BitTorrent"
Direction     Status                    Granted                  Used     % Used
-------------------------------------------------------------------------------
to sub        valid                     1378168              137816        10%

from sub      valid                     1381148                13781         1%
both          valid        18446744073709551614                151597         5%
===============================================================================
*A:BNG-1>show>app-assure>group#
```

The PCRF can also at the same time set ADC or PCC rules for other applications (such as the *0_rated* and the default_charging_group charging groups).

In the following case, the PCRF installs an ADC usage monitoring rule for:

- Charging group: “0-rated”, but without usage thresholds
- Charging group: “default_charge_group”, and sets only a threshold for “to sub” traffic.

This results in having a usage policy for the “0-rated” charging group installed but this is not active since there are no grants associated with it:

```
*A:BNG-1>show>app-assure>group# aa-sub esm "alcatel_A_1" usage-monitor status
===============================================================================
Application-Assurance Subscriber "alcatel_A_1" (esm)
Usage Monitor Status
===============================================================================
Type            Name                              Oper Status
```

```
-------------------------------------------------------------------------------
application      BitTorrent                      active
charging-group   0_rated                         inactive
charging-group   default_charge_group            active
-------------------------------------------------------------------------------
No. of rules: 3
===============================================================================
*A:BNG-1>show>app-assure>group#
```

Note that the “inactive” status for the “0-rated” charging group is due to no grants being received.

Moreover, detailed counters show:

```
*A:BNG-1>show>app-assure>group# aa-sub esm "alcatel_A_1" usage-monitor count
===============================================================================
Application-Assurance Subscriber "alcatel_A_1" (esm)
Usage Monitor Credit Statistics
===============================================================================
Application: "BitTorrent"
Direction     Status                    Granted                 Used     % Used
-------------------------------------------------------------------------------
to sub        valid                     1378168             137816        10%

from sub      valid                     1381148                13781         1%
both          valid        18446744073709551614               151597         5%
-------------------------------------------------------------------------------
Charging-Group: "0_rated"
Direction     Status                    Granted                 Used     % Used
-------------------------------------------------------------------------------
to sub        invalid                       n/a                    0        n/a
from sub      invalid                       n/a                    0        n/a
both          invalid                       n/a                    0        n/a
-------------------------------------------------------------------------------
Charging-Group: "default_charge_group"
Direction     Status                    Granted                 Used     % Used
-------------------------------------------------------------------------------
to sub        valid                     1000000              1378084       100%
from sub      invalid                       n/a                 1574        n/a
both          invalid                       n/a              1379658        n/a
===============================================================================
*A:BNG-1>show>app-assure>group#
```

Again, the “invalid” status above reflects the fact that no grants have been received.

Conclusion

The introduction of the diameter (/Gx) control feature on the 7x50 BNG enables operators to consolidate policy management systems used in wire-line and wireless environments into a single system. This provides an increase in operational efficiency as mobile and fixed networks convergence gains more traction.

This example illustrates how policy control and usage monitoring of the 7x50 BNG Application Assurance services can be achieved over standard 3GPP Diameter Gx protocol.

Deterministic Large Scale NAT44

This chapter provides information about deterministic large scale NAT44 configurations.

Topics in this chapter include:

Applicability

The information and configuration in this chapter was based on release 11.0.R3, and is updated to release 14.0.R4.

Overview

Deterministic Network Address Translation (NAT) is a mode of operation where mappings between the NAT subscriber and the outside IP address and port range are allocated at the time of configuration.

In deterministic NAT for Large Scale NAT IPv4-to-IPv4 (LSN44) subscribers, each LSN44 subscriber is permanently mapped to an outside IP address and a dedicated (deterministic) port-block based on a specific algorithm.

Logging is not needed in this case because the reverse mapping can be obtained using the reverse of the preceding algorithm.

A deterministic LSN44 subscriber can have only one deterministic port-block that can (optionally) be extended by one or multiple dynamic port-blocks in case all ports in deterministic port-block are exhausted.

In case an LSN44 subscriber has been assigned both deterministic and dynamic port blocks, logging for the dynamic port-block allocation/de-allocation is required.

A scalable logging solution for dynamic port-blocks is achievable using RADIUS or IPFIX.

Logging for dynamic port-blocks is out of the scope of this chapter.

Figure 40 **Deterministic NAT Mapping**

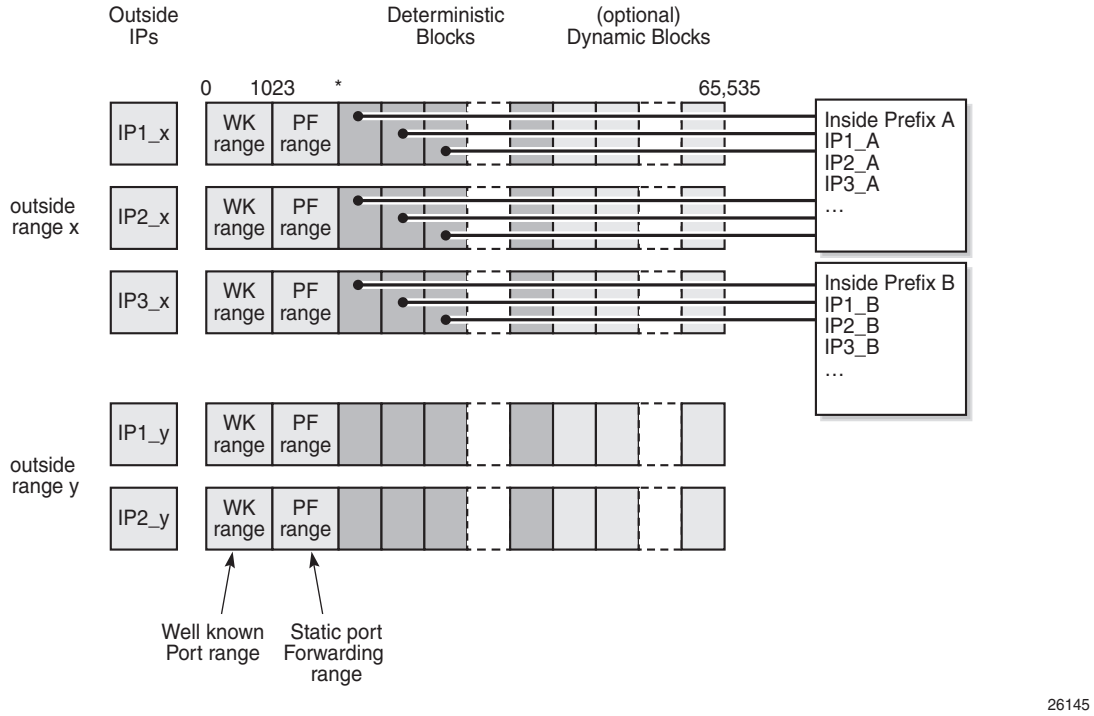

Algorithm

The deterministic NAT algorithm makes a predictable mapping between the (inside IP, routing instance) and the (outside IP, routing instance, deterministic port block).

The algorithm is revertive, meaning that a given (outside IP, routing instance, deterministic port block) will derive a given (inside IP, routing Instance).

The algorithm is loosely based on the draft RFC draft-donley-behave-deterministic-cgn-00.txt, which allows for the dynamic expansion of the port-blocks once the ports in the original deterministic port-block are exhausted.

Figure 41 **Deterministic NAT Algorithm**

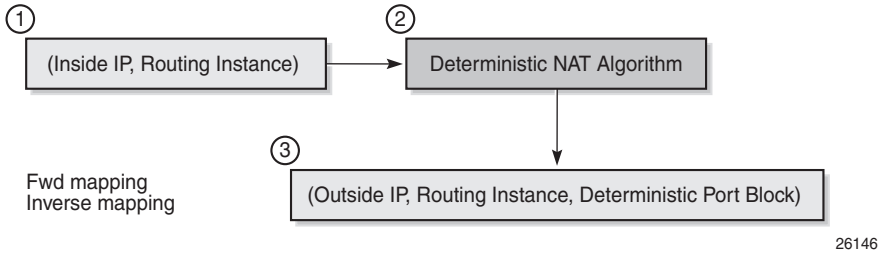

Deterministic Mapping

Any inside prefix in any routing instance can be mapped to any pool in any routing instance.

In deterministic NAT, prefixes from multiple routing instances can be mapped to the same outside pool, also prefixes from a single inside routing instance can be selectively mapped to different outside pools.

Figure 42 **Deterministic Mapping: Inside -> Outside Routing Instances**

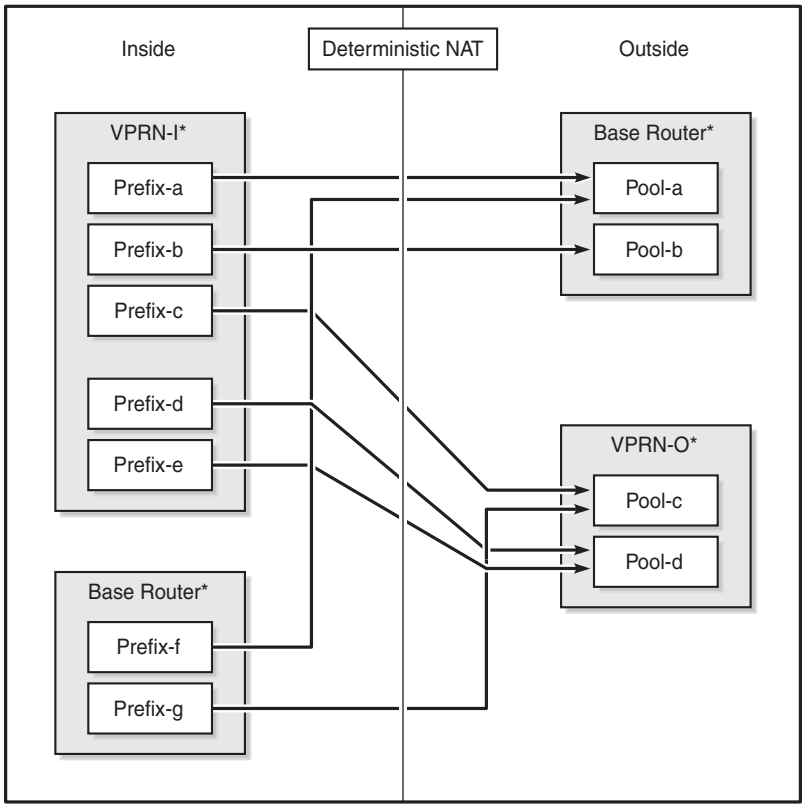

*Routing-Based NAT cannot be used if inside/outside routing instances are the same

26147

Mapping Rules

A deterministic LSN44 subscriber is mapped to only one deterministic block which can further be extended to multiple dynamic blocks if ports within the deterministic block are exhausted.

The subscriber-limit is the number of subscribers that can be deterministically mapped to an outside IP address (i.e. compression ratio) and *must* be a power of 2.

The total number of deterministic ports (DetP) per outside IP address is determined by the number of subscribers per outside IP address and the number of deterministic ports per subscriber.

The remaining ports (DynP) beyond the deterministic port range up to 65535 will be dedicated for dynamic use when a deterministic block is exhausted.

Every host using an inside prefix is guaranteed one dedicate block in the deterministic port ranges.

If the inside prefix length is m < 32-n, where 2^n=subscriber-limit, then the prefix must be broken into pieces so that all hosts (subscriber-limit) in each piece maps exactly to one outside IP address.

For example, if there is an inside prefix 192.168.0.0/23, here m=23; and the subscriber-limit is also set to 256, then n=8. This results in 23 < 24 (32-8) and so this inside prefix has to be broken into 2 pieces, in other words, this inside prefix will fit into 2 outside IP addresses, each of 256 port-blocks.

In case that the prefix length is m ≥ 32-n, where 2^n=subscriber-limit, then all hosts from the configured prefix are mapped to the same outside IP.

For example, if there is an inside prefix 192.168.1.0/25, here m=25, and there can be at most 128 hosts, and the subscriber-limit is set to 256, then n=8. This results in 25 > 24 (32-8), so definitely 128 hosts can fit in one outside IP because there are 256 available port-blocks, in other words, this inside prefix will fit into one outside IP where 128 blocks have been used out of the 256 port-blocks available, and the rest (256-128) are wasted.

Overbooking of the outside address pool is not supported in deterministic NAT.

Figure 43 **Deterministic Mapping: Outside IP Port-Blocks/Ranges**

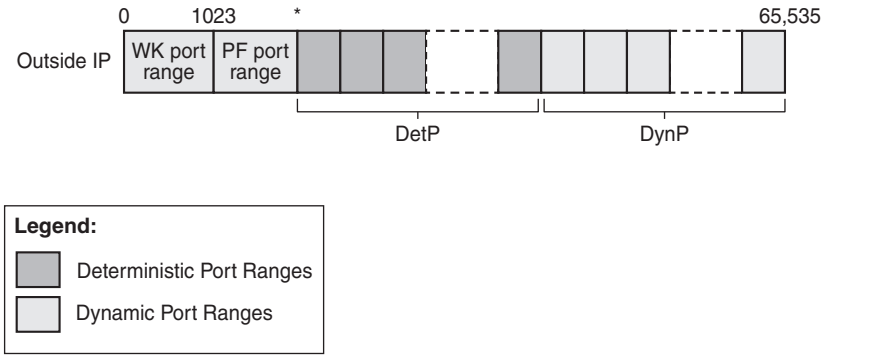

Configuration

Figure 44 **Example Topology**

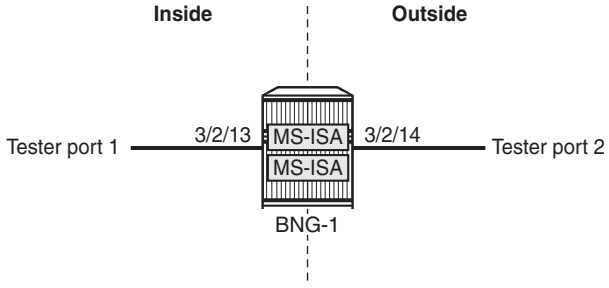

Configuration Prerequisites

Chassis mode, card, and MDA configuration.

```
configure
    system
        chassis-mode c
    exit
exit

configure
    card 2
        card-type iom3-xp
        mda 1
            mda-type isa-bb
            no shutdown
        exit
        mda 2
            mda-type isa-bb
            no shutdown
        exit
        no shutdown
    exit
exit
```

Note: Private address ranges are used in outside pools within this chapter but normally public address ranges would be used.

Create the nat-group, and add the MS-ISAs created above to the nat-group; up to 10 MS-ISAs of type isa-bb can be configured under the nat-group.

```
configure
    isa
        nat-group 1 create
            mda 2/1
            mda 2/2
            active-mda-limit 1
            no shutdown
        exit
    exit
exit
```

Configuration Commands

A NAT **outside pool** is configured using the following command:

```
configure  {router | service vprn <service-id>}
  nat
    outside
      pool <nat-pool-name> [nat-group <nat-group-id> type <pool-type> create]
        port-reservation {blocks <num-blocks> | ports <num-ports>}
        port-forwarding-range <range-end>
        subscriber-limit <subscriber-limit>
        deterministic
        port-reservation <num-ports>
        exit
          address-range <start-ip-address> <end-ip-address> create
        exit
      exit
    exit
  exit
```

where:

nat-pool-name — Specifies the name of the NAT pool up to 32 characters max.

nat-group-id — Specifies the NAT group ID. The values are 1 — 4.

pool-type — Species the pool type (**large-scale**).

num-blocks — Specifies the number of port-blocks per IP address. Setting num-blocks to one (1) for large scale NAT will enable 1:1 NAT for IP addresses in this pool The values are 1 — 64512

num-ports — Specifies the number of ports per block. The values are 1 — 32256

range-end — Specifies the end of the port range available for port forwarding. The values are 1023 — 65535

subscriber-limit — Specifies the maximum number of subscribers per IP address. A power of 2 (2^n) number for deterministic NAT

[1,2,4,8,16,32,64,128,256,512,1024,2048, 4096, 8192,16348, 32768]
1..65535 for non-deterministic NAT
default: 65535 for non-deterministic

num-ports — Specifies the number of ports in a deterministic port block that is allocated and dedicated to a single subscribers during the configuration phase. The values are 1..65535

start-ip-address — Specifies the beginning IP address in the a.b.c.d format.

end-ip-address — Specifies the ending IP address in the a.b.c.d. format.

Note:

- When the subscriber-limit equals 1, each subscriber is mapped to a single outside IP address, though the NAPT (port translation) function is still performed.
- 1:1 NAT mode in combination with deterministic NAT is not supported.

A NAT **policy** is configured using the following command:

```
configure service nat
 nat-policy <nat-policy-name> [create]
    block-limit <[1..40]>
    pool <nat-pool-name> {router <router-instance> | service-name <service-name>}
 exit
```

where:

nat-policy-name — Specifies the NAT policy name up to 32 characters max.

block-limit —The maximum number of deterministic plus dynamic port blocks that can be assigned to a single inside IP address. In other words, the maximum number of dynamic port blocks that can be assigned to an inside IP address when the deterministic port block is exhausted equals (block-limit - 1).

nat-pool-name — Specifies the NAT pool name up to 32 characters max.

router-instance — Specifies the router instance the pool belongs to, either by router name or service ID.

<router-name>|<service-id>

The router name values are **Base** or *service-id* [1..2147483647]

service-name — Specifies the name of the service up to 64 characters max.

A NAT **inside prefix** is configured using the following command:

```
configure [router| service vprn <service-id>]
  nat
    inside
      deterministic
        classic-lsn-max-subscriber-limit <max>
        prefix <ip-prefix/length> subscriber-type <nat-sub-type>
               nat-policy <nat-policy-name> create
          map start <lsn-sub-address> end <lsn-sub-address> to <outside-ip-address>
          no shutdown
          exit
        exit
      exit
    exit
```

where:

max — The power of 2 (2^n) number that must match the largest subscriber limit number in a deterministic pool referenced from this inside routing instance. The range for this command is the same as the subscriber-limit command under the pool hierarchy. The values are 1,2,4,8 — 32768

ip-prefix/length — A prefix on the inside encompassing subscribers that will be deterministically mapped to an outside IP address and port block in the corresponding pool.

<*ip-prefix/ip-pref**>	<ipv4-prefix>/<ipv4-prefix-length> |
	<ipv6-prefix>/<ipv6-prefix-length>
<*ipv4-prefix*>	a.b.c.d (host bits must be 0)
<*ipv4-prefix-length*>	[0..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..128]
<*nat-sub-type*>:	classic-lsn-sub
<*nat-policy-name*>	Specifies a NAT policy name up to 32 characters in length.

Following rules apply to the *classic-lsn-max-subscriber-limit*:

- Should be greater than the largest subscriber-limit of all pools referenced by the NAT policies within the corresponding inside routing instance.
- Must be configured before any inside prefix configuration.
- Must be 2^n and affects the ingress hashing of deterministic subscribers and also non-deterministic subscribers in case both are configured under the same inside router instance.

Three cases are now configured to demonstrate the use of deterministic and dynamic port-block usage:

- Case 1: Mapping multiple prefixes from the same VRF into the same outside pool.
- Case 2: Mapping multiple prefixes from the same VRF into different outside pools.
- Case 3: Mapping overlapping prefixes from different VRFs into the same outside pool.

In each case all of the traffic is NATed.

Case 1

Configured with:

- Mapping multiple prefixes of the same VRF into the same outside pool.
- NAT all traffic.

Figure 45 **Case 1**

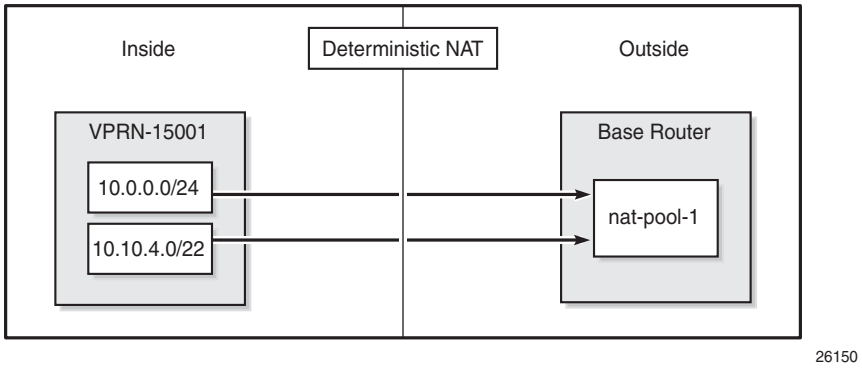

The NAT **outside pool** is configured as follows:

```
configure
    router
        nat
            outside
                pool "nat-pool-1" nat-group 1 type large-scale create
                    port-reservation ports 180
                    port-forwarding-range 4023
                    subscriber-limit 128
                    deterministic
                        port-reservation 300
```

```
                    exit
                    address-range 192.168.0.1 192.168.0.100 create
                    exit
                    no shutdown
                exit
            exit
        exit
    exit
exit
```

The NAT **policy** is configured as follows:

```
configure
    service
        nat
            nat-policy "nat-policy-1" create
                block-limit 4
                pool "nat-pool-1" router Base
            exit
        exit
    exit
exit
```

The NAT **inside prefix** is configured as follows:

```
configure
    service
        vprn 15001 customer 1 create
            nat
                inside
                    destination-prefix 0.0.0.0/0
                    classic-lsn-max-subscriber-limit 256
                    deterministic
                        prefix 10.0.0.0/24 subscriber-type classic-lsn-sub
                                                nat-policy "nat-policy-1" create
                            map start 10.0.0.0 end 10.0.0.255 to 192.168.0.1
                            no shutdown
                        exit

                        prefix 10.10.4.0/22 subscriber-type classic-lsn-sub
                                                 nat-policy "nat-policy-1" create
                            map start 10.10.4.0 end 10.10.7.255 to 192.168.0.3
                            no shutdown
                        exit
                    exit
                exit
            exit
            no shutdown
        exit
    exit
exit
```

The **classic-lsn-max-subscriber-limit** value should be greater or equal to the largest subscriber-limit of all pools referenced by NAT policies within the corresponding inside routing instance. It must be 2^n and affects ingress hashing of deterministic subscribers.

map statements are automatically created when the prefix is created and it is **no shutdown**.

Show Commands

The subscriber-limit is set to 128 for the 10.0.0.0/24 prefix, so it is broken into two smaller /25 prefixes each. Each of these smaller prefixes are mapped into a specific outside IP address.

To show the first Large Scale NAT (LSN) subscriber of the first /25 prefix for inside routing instance 15001, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-router 15001 inside-ip 10.0.0.0

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.0.0.0]
NAT policy                  : nat-policy-1
Subscriber ID               : 276824064
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.0.0.0/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : "Base"
Outside IP address          : 192.168.0.1

-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
*A:PE1#
```

The last subscriber mapping to the same 192.168.0.1 outside IP address has inside address 10.0.0.127.

To show the first LSN subscriber of the second /25 prefix for inside routing instance 15001, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-router 15001 inside-ip 10.0.0.128

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.0.0.128]
NAT policy                  : nat-policy-1
Subscriber ID               : 276824192
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.0.0.128/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : "Base"
Outside IP address          : 192.168.0.2

-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
*A:PE1#
```

The last subscriber mapping to the same 192.168.0.2 outside IP address has inside address 10.0.0.255.

To show the base router LSN blocks corresponding to the first inside IP address within the 10.0.0.0/24 prefix, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.0.0.0

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.1 [4024..4323]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2016/10/27 11:18:59
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.0

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
*A:PE1#
```

To show the base router LSN blocks corresponding to the last inside IP address within the 10.0.0.0/24 prefix, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.0.0.255

===============================================================================
Large-Scale NAT blocks for Base
```

```
===============================================================================
192.168.0.2 [42124..42423]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2016/10/27 11:18:59
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.255

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
*A:PE1#
```

The subscriber-limit is 128 for the 10.10.4.0/22 prefix, so it is broken into eight /25 prefixes. Each of these smaller prefixes are mapped into a specific outside IP address.

To show the first LSN subscriber of the first /25 prefix for inside routing instance 15001, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-ip 10.10.4.0

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.10.4.0]
NAT policy                  : nat-policy-1
Subscriber ID               : 276824320
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.10.4.0/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : "Base"
Outside IP address          : 192.168.0.3

-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
*A:PE1#
```

The last subscriber mapping to the same 192.168.0.3 outside IP address has inside address 10.10.4.127.

To show the first LSN subscriber of the second /25 prefix for inside routing instance 15001, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-ip 10.10.7.128

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.10.7.128]
NAT policy                  : nat-policy-1
```

```
Subscriber ID               : 276825216
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.10.7.128/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : "Base"
Outside IP address          : 192.168.0.10

-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
*A:PE1#
```

To show the base router LSN blocks corresponding to the first inside IP within 10.10.4.0/24 prefix, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.10.4.0

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.3 [4024..4323]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2016/10/27 11:18:59
Inside router                       : vprn15001
Inside IP address                   : 10.10.4.0

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
*A:PE1#
```

To show the base router LSN blocks corresponding to the last inside IP within 10.10.4.0/24 prefix, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.10.7.255

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.10 [42124..42423]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2016/10/27 11:18:59
Inside router                       : vprn15001
Inside IP address                   : 10.10.7.255

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
*A:PE1#
```

Mapping Results

According to this configuration, each inside IP address has one deterministic block of 300 ports and can have up to three dynamic blocks (block-limit = 4) each of 180 ports, allowing a maximum of 300+3*180 = 840 flows.

Figure 46 **Case 1 Results**

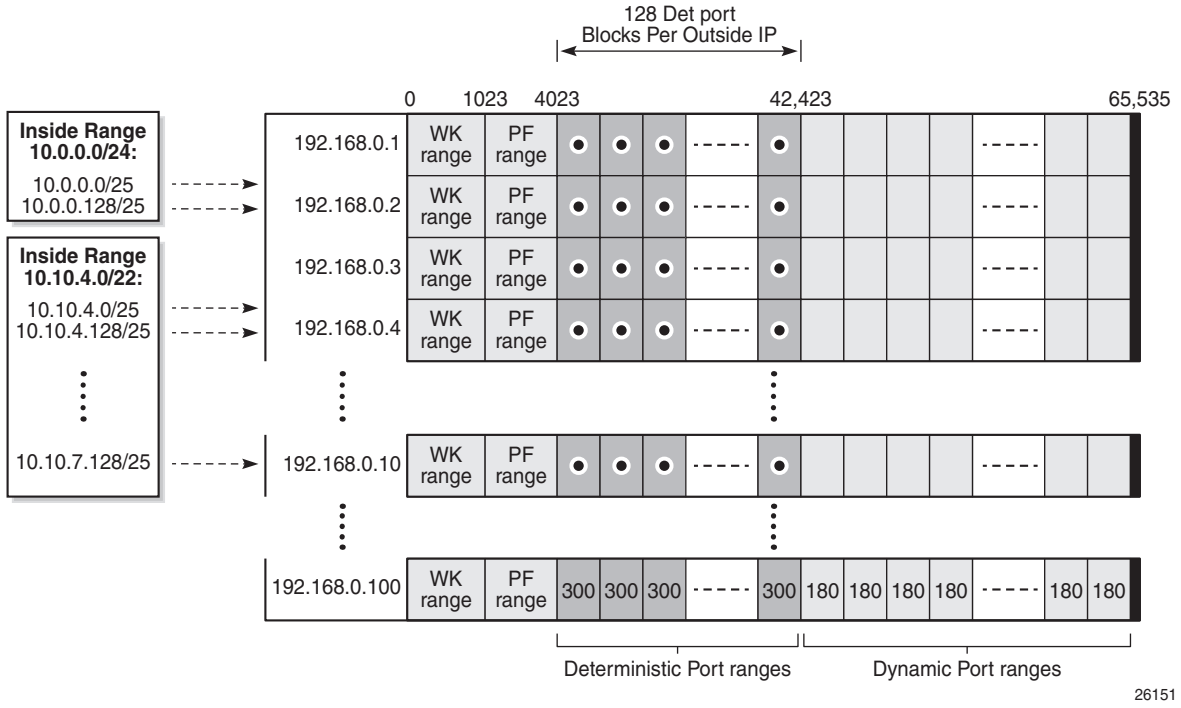

Sending Flows

For the inside IP 10.0.0.1, several UDP flows will be sent and both the deterministic and dynamic blocks mappings will be verified.

Figure 47 **Case 1 Flows**

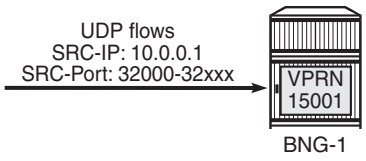

When sending 300 UDP flows or less, all flows are mapped to a single deterministic block because the number of ports in a deterministic block is 300. There is no logging; because no dynamic blocks are used, and only the deterministic block is used.

To show LSN blocks on the outside routing instance **Base** and the outside ports allocated for the inside IP 10.0.0.1, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.0.0.1

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.1 [4324..4623]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2016/10/27 11:18:59
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
*A:PE1#
```

When increasing the number of flows such that: 301 ≤ number of flows ≤ 480

- In addition to the deterministic block (300 ports), there will be an extension by 1 dynamic block of 180 ports (port-reservation=180).
- Logging occurs for the dynamic port-block.

To show the base router LSN blocks and the outside ports allocated to the inside IP address 10.0.0.1, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.0.0.1

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.1 [4324..4623]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2016/10/27 11:18:59
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

192.168.0.1 [44044..44223]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2016/10/28 12:40:41
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

-------------------------------------------------------------------------------
```

```
Number of blocks: 2
===============================================================================
*A:PE1#
```

Logging is verified using Log 99 (in case event-control **nat** events are generated) which shows the mapping details to the new dynamic block as follows:

```
2 2016/10/28 12:40:41.51 UTC MINOR: NAT #2012 Base NAT
"{12} Map  192.168.0.1 [44044-44223] MDA 2/1 -- 276824065 classic-lsn-
sub %1 vprn15001 10.0.0.1 at 2016/10/28 12:40:41"
```

When increasing the number of flows such that: 481 ≤ number of flows ≤ 660

- In addition to the deterministic block (300 ports), there will be an extension by 2 dynamic blocks of 180 ports each.
- Logging occurs for the dynamic port-blocks.

To show LSN blocks on the outside routing instance **Base** and the outside ports allocated for the inside IP 10.0.0.1, the following command is used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.0.0.1

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.1 [4324..4623]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2016/10/27 11:18:59
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

192.168.0.1 [44044..44223]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2016/10/28 12:40:41
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

192.168.0.1 [44224..44403]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2016/10/28 12:41:52
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

-------------------------------------------------------------------------------
Number of blocks: 3
===============================================================================
*A:PE1#
```

Logging is verified using Log 99 (in case event-control **nat** events are generated) which shows the mapping details to the new dynamic block as follows:

```
3 2016/10/28 12:41:52.66 UTC MINOR: NAT #2012 Base NAT
"{13} Map  192.168.0.1 [44224-44403] MDA 2/1 -- 276824065 classic-lsn-
sub %1 vprn15001 10.0.0.1 at 2016/10/28 12:41:52"
```

When increasing the number of flows such that :661 ≤ number of flows ≤ 840

- In addition to the deterministic block (300 ports), there will be an extension by 3 dynamic blocks of 180 ports each.
- Logging occurs for the dynamic port-blocks.

To show LSN blocks on the outside routing instance “Base” and the outside ports allocated for the inside IP 10.0.0.1, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.0.0.1

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.0.1 [4324..4623]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2016/10/27 11:18:59
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

192.168.0.1 [44044..44223]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2016/10/28 12:40:41
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

192.168.0.1 [44224..44403]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2016/10/28 12:41:52
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

192.168.0.1 [44404..44583]
Pool                                : nat-pool-1
Policy                              : nat-policy-1
Started                             : 2016/10/28 12:43:46
Inside router                       : vprn15001
Inside IP address                   : 10.0.0.1

-------------------------------------------------------------------------------
Number of blocks: 4
===============================================================================
*A:PE1#
```

Logging is verified using Log 99 (in case event-control **nat** events are generated) which shows the mapping details to the new dynamic block as follows:

```
4 2016/10/28 12:43:46.71 UTC MINOR: NAT #2012 Base NAT
```

```
"{14} Map  192.168.0.1 [44404-44583] MDA 2/1 -- 276824065 classic-lsn-
sub %1 vprn15001 10.0.0.1 at 2016/10/28 12:43:46"
```

When increasing number of flows such that the number of flows > 840

- No more extension by dynamic blocks (block-limit = 4) allowed.
- Any flows more than 840 will be dropped and the relevant NAT statistics incremented.

To verify NAT statistics, first check the NAT group/member and MS-ISA associated with the outside IP 192.168.0.1/32:

```
*A:PE1# show router route-table 192.168.0.1/32

===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
192.168.0.1/32                                  Remote  NAT       01d01h26m  0
       NAT outside to mda 2/1                                          0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
*A:PE1#
```

To check which group/member does this MS-ISA belong to, the following command can be used:

```
*A:PE1# show isa nat-group 1 members

===============================================================================
ISA Group 1 members
===============================================================================
Group Member     State          Mda  Addresses  Blocks     Se-% Hi Se-Prio
-------------------------------------------------------------------------------
1     1          active         2/1  175        23648      < 1  N  0
-------------------------------------------------------------------------------
No. of members: 1
===============================================================================
*A:PE1#
```

To verify relevant statistics for this NAT group/member, the following command can be used:

```
*A:PE1# show isa nat-group 1 member 1 statistics | match flow
no matching flow                                  : 56818
max flow exceeded                                 : 0
TCP no flow for RST                               : 0
```

```
TCP no flow for FIN                                    : 0
TCP no flow                                            : 128094
flow log failed                                        : 0
new flow                                               : 1470768
found flow                                             : 39661850
flow create logged                                     : 0
flow delete logged                                     : 0
flow log pkt tx                                        : 0
flow create failed, key ambiguous                      : 0
flow create failed, conflicting policies               : 0
*A:PE1#
```

Case 2

Configured with:

- Mapping multiple prefixes from the same VRF into different outside pools.
- NAT all traffic.

Figure 48 **Case 2**

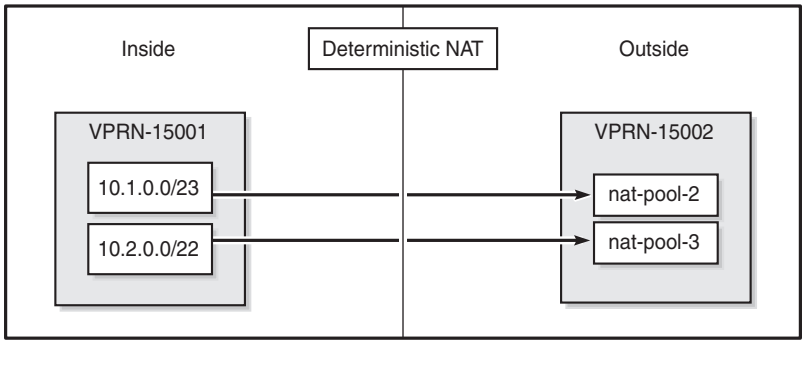

The NAT **outside pools** are configured as follows:

```
configure
    service
        vprn 15002 customer 1 create
            nat
                outside
                    pool "nat-pool-2" nat-group 1 type large-scale create
                        port-reservation ports 80
                        subscriber-limit 256
                        deterministic
                            port-reservation 180
                        exit
                        address-range 192.168.2.1 192.168.2.200 create
                        exit
                        no shutdown
```

```
                    exit
                    pool "nat-pool-3" nat-group 1 type large-scale create
                        port-reservation ports 120
                        port-forwarding-range 4023
                        subscriber-limit 64
                        deterministic
                            port-reservation 840
                        exit
                        address-range 192.168.3.1 192.168.3.200 create
                        exit
                        no shutdown
                    exit
                exit
            exit
        exit
    exit
exit
```

The NAT **policies** are configured as follows:

```
configure
    service
        nat
            nat-policy "nat-policy-2" create
                block-limit 4
                pool "nat-pool-2" router 15002
            exit
            nat-policy "nat-policy-3" create
                block-limit 2
                pool "nat-pool-3" router 15002
            exit
        exit
    exit
exit
```

The NAT **inside prefix** is configured as follows:

```
configure
    service
        vprn 15001 customer 1 create
            nat
                inside
                    destination-prefix 0.0.0.0/0
                    classic-lsn-max-subscriber-limit 256
                    deterministic
                        prefix 10.1.0.0/23 subscriber-type classic-lsn-sub
                                                nat-policy "nat-policy-2" create
                            map start 10.1.0.0 end 10.1.1.255 to 192.168.2.1
                            no shutdown
                        exit
                        prefix 10.2.0.0/22 subscriber-type classic-lsn-sub
                                                nat-policy "nat-policy-3" create
                            map start 10.2.0.0 end 10.2.3.255 to 192.168.3.1
                            no shutdown
                        exit
                    exit
                exit
```

```
            exit
        exit
    exit
exit
```

Show Commands

The subscriber-limit corresponding to the 10.1.0.0/23 prefix is 256, so the 10.1.0.0/23 prefix is broken into two /24 prefixes. Each of these smaller prefixes are mapped into a specific outside IP address.

To show the first LSN subscriber of the first /24 prefix for inside routing instance 15001, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-ip 10.1.0.0

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.1.0.0]
NAT policy                  : nat-policy-2
Subscriber ID               : 276829472
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.1.0.0/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : 15002
Outside IP address          : 192.168.2.1

-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
*A:PE1#
```

The last subscriber mapping to the same 192.168.2.1 outside IP address has inside address 10.1.0.255.

To show the first LSN subscriber of the second /24 prefix for inside routing instance 15001, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-ip 10.1.1.0

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.1.1.0]
NAT policy                  : nat-policy-2
Subscriber ID               : 276829728
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
```

```
Inside IP address prefix    : 10.1.1.0/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : 15002
Outside IP address          : 192.168.2.2

-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
*A:PE1#
```

The last subscriber mapping to the same 192.168.2.2 outside IP address has inside address 10.1.1.255.

To show the VPRN-15002 LSN blocks corresponding to the first inside IP address within 10.1.0.0/23 prefix, the following command can be used:

```
*A:PE1# show router 15002 nat lsn-blocks inside-ip 10.1.0.0

===============================================================================
Large-Scale NAT blocks for vprn15002
===============================================================================
192.168.2.1 [1024..1203]
Pool                                : nat-pool-2
Policy                              : nat-policy-2
Started                             : 2016/10/28 12:53:23
Inside router                       : vprn15001
Inside IP address                   : 10.1.0.0

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
*A:PE1#
```

To show the VPRN-15002 LSN blocks corresponding to the last inside IP address within 10.1.0.0/23 prefix, the following command can be used:

```
*A:PE1# show router 15002 nat lsn-blocks inside-ip 10.1.1.255

===============================================================================
Large-Scale NAT blocks for vprn15002
===============================================================================
192.168.2.2 [46924..47103]
Pool                                : nat-pool-2
Policy                              : nat-policy-2
Started                             : 2016/10/28 12:53:23
Inside router                       : vprn15001
Inside IP address                   : 10.1.1.255

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
*A:PE1#
```

The subscriber-limit corresponding to the 10.2.0.0/22 prefix is sixty four,so the 10.2.0.0/22 prefix is broken into sixteen /26 prefixes. Each of these /26 prefixes is mapped to a specific outside IP address.

To show the first LSN subscriber for the inside routing instance 15001 for the first /26 prefix, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-ip 10.2.0.0

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.2.0.0]
NAT policy                  : nat-policy-3
Subscriber ID               : 276829984
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.2.0.0/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : 15002
Outside IP address          : 192.168.3.1

-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
*A:PE1#
```

The last inside address mapping to the 192.168.3.1 outside address is 10.2.0.63.

To show the first LSN subscriber for the inside routing instance 15001 for the last /26 prefix, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-ip 10.2.3.192

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.2.3.192]
NAT policy                  : nat-policy-3
Subscriber ID               : 276830944
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.2.3.192/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : 15002
Outside IP address          : 192.168.3.16

-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
*A:PE1#
```

The last inside address mapping to the 192.168.3.16 outside address is 10.2.3.255.

To show the VPRN-15002 LSN blocks corresponding to the first inside IP address within the 10.2.0.0/22 prefix, the following command can be used:

```
*A:PE1# show router 15002 nat lsn-blocks inside-ip 10.2.0.0

===============================================================================
Large-Scale NAT blocks for vprn15002
===============================================================================
192.168.3.1 [4024..4863]
Pool                                : nat-pool-3
Policy                              : nat-policy-3
Started                             : 2016/10/28 12:53:23
Inside router                       : vprn15001
Inside IP address                   : 10.2.0.0

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
*A:PE1#
```

To show the VPRN-15002 LSN blocks corresponding to the last inside IP within 10.2.0.0/22 prefix, the following command can be used:

```
*A:PE1# show router 15002 nat lsn-blocks inside-ip 10.2.3.255

===============================================================================
Large-Scale NAT blocks for vprn15002
===============================================================================
192.168.3.16 [56944..57783]
Pool                                : nat-pool-3
Policy                              : nat-policy-3
Started                             : 2016/10/28 12:53:23
Inside router                       : vprn15001
Inside IP address                   : 10.2.3.255

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
*A:PE1#
```

Mapping results

According to this configuration, for the 10.1.0.0/23 prefix, each inside IP address has one deterministic block of 180 ports and can have up to three dynamic blocks (block-limit =4) each of 80 ports, allowing for a maximum of 180+3*80 = 420 flows.

Figure 49 **Case 2: Prefix 10.1.0.0/23 Results**

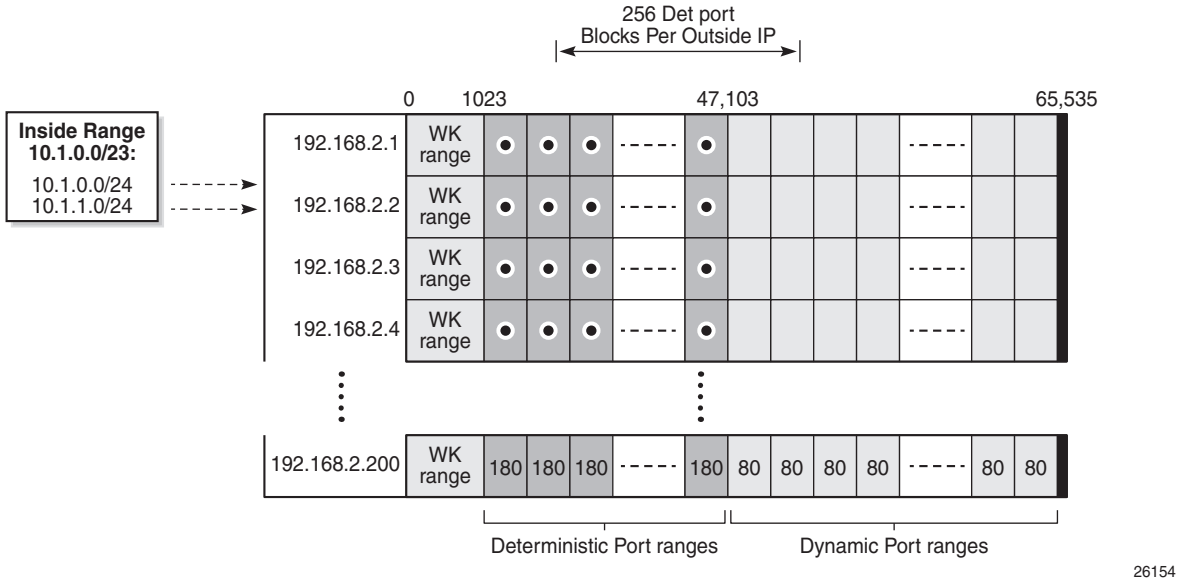

According to this configuration, for the 10.2.0.0/22 prefix, each inside IP address has one deterministic block of 840 ports, and can have up to one dynamic block (block-limit =2) of 120 ports, allowing for a maximum of 840+120 = 960 flows.

Figure 50 **Case 2: Prefix 10.2.0.0/22 Results**

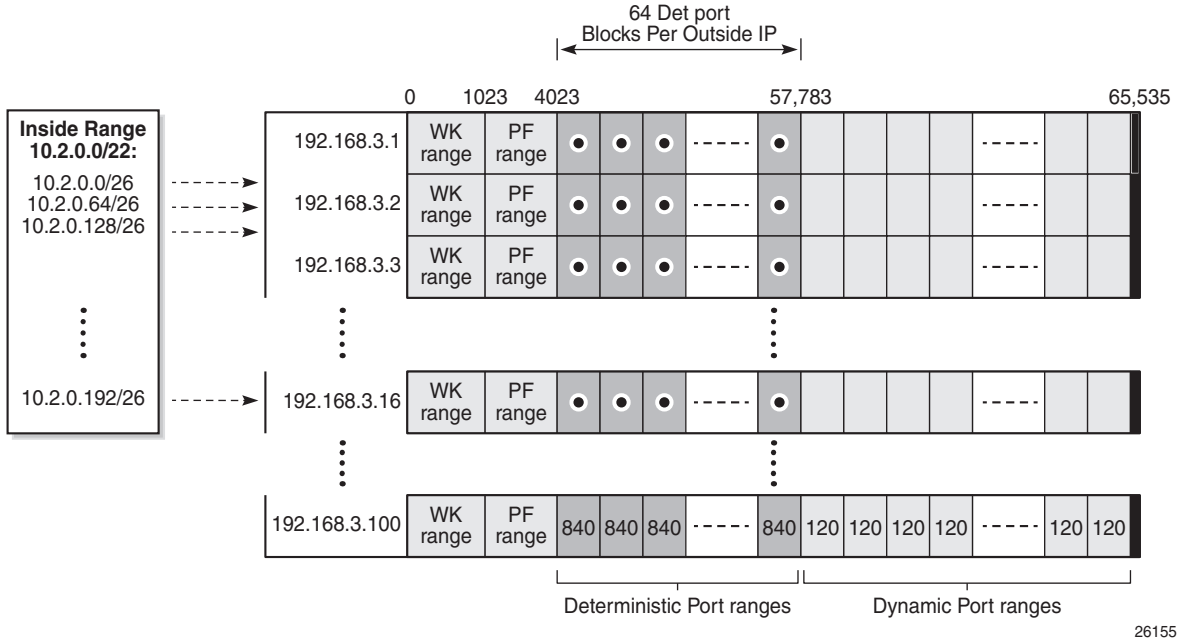

Case 3

Configured with:

- Mapping overlapping prefixes from different VRFs into the same outside pool.
- NAT all traffic.

Figure 51 **Case 3**

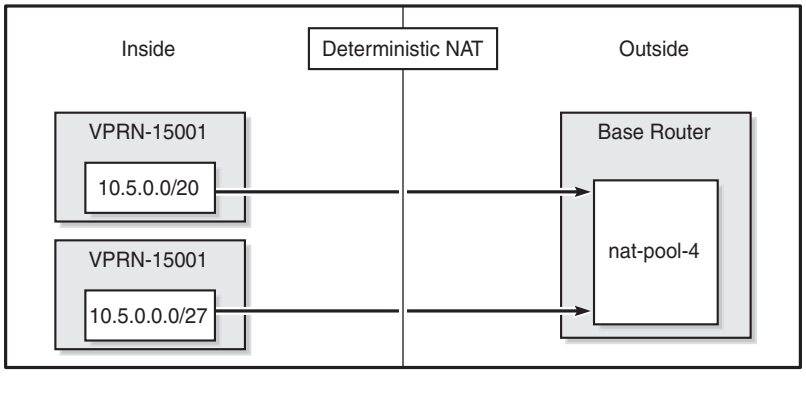

The NAT **outside pool** is configured as follows:

```
configure
    router "Base"
        nat
            outside
                pool "nat-pool-4" nat-group 1 type large-scale create
                    port-reservation ports 461
                    port-forwarding-range 4023
                    subscriber-limit 64
                    deterministic
                        port-reservation 500
                    exit
                    address-range 192.168.4.1 192.168.4.100 create
                    exit
                    no shutdown
                exit
            exit
        exit
    exit
exit
```

The NAT **policy** is configured as follows:

```
configure
    service
        nat
            nat-policy "nat-policy-4" create
```

```
                block-limit 4
                pool "nat-pool-4" router Base
            exit
        exit
    exit
exit
```

The NAT **inside prefix** is configured as follows:

```
configure
    service
        vprn 15001 customer 1 create
            nat
                inside
                    destination-prefix 0.0.0.0/0
                    classic-lsn-max-subscriber-limit 256
                    deterministic
                        prefix 10.5.0.0/20 subscriber-type classic-lsn-sub
                                                nat-policy "nat-policy-4" create
                            map start 10.5.0.0 end 10.5.15.255 to 192.168.4.1
                            no shutdown
                        exit
                    exit
                exit
            exit
        exit
    exit
exit

configure
    service
        vprn 15002 customer 1 create
            nat
                inside
                    destination-prefix 0.0.0.0/0
                    classic-lsn-max-subscriber-limit 128
                    deterministic
                        prefix 10.5.0.0/27 subscriber-type classic-lsn-sub
                                                nat-policy "nat-policy-4" create
                            map start 10.5.0.0 end 10.5.0.31 to 192.168.4.65
                            no shutdown
                        exit
                    exit
                exit
            exit
        exit
    exit
exit
```

Show Commands

For the 10.5.0.0/20 prefix on VPRN 15001, the subscriber-limit is 64.The 10.5.0.0/20 prefix will be broken into 64 smaller /26 prefixes, each will be mapped into a specific outside IP address.

To show the first LSN subscriber for the inside routing instance 15001 of the first /26 prefix, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-ip 10.5.0.0 inside-router 15001

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.5.0.0]
NAT policy                  : nat-policy-4
Subscriber ID               : 276825344
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.5.0.0/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : "Base"
Outside IP address          : 192.168.4.1

-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
*A:PE1#
```

The last inside address mapping to the 192.168.4.1 outside address is 10.5.0.63.

To show the first Large Scale NAT (LSN) subscriber for the inside routing instance 15001 of the last /26 prefix, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-ip 10.5.15.192 inside-router 15001

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.5.15.192]
NAT policy                  : nat-policy-4
Subscriber ID               : 276829376
-------------------------------------------------------------------------------
Type                        : classic-lsn-sub
Inside router               : 15001
Inside IP address prefix    : 10.5.15.192/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : "Base"
Outside IP address          : 192.168.4.64

-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
*A:PE1#
```

The last inside address mapping to the 192.168.4.64 outside address is 10.5.15.255.

To show the base router LSN blocks corresponding to the first inside IP address within the 10.5.0.0/20 prefix, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.5.0.0 inside-router 15001

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.4.1 [4024..4523]
Pool                                  : nat-pool-4
Policy                                : nat-policy-4
Started                               : 2016/10/27 13:11:38
Inside router                         : vprn15001
Inside IP address                     : 10.5.0.0

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
*A:PE1#
```

To show the base router LSN blocks corresponding to the last inside IP address within the 10.5.0.0/20 prefix, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.5.15.255 inside-router 15001

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.4.64 [35524..36023]
Pool                                  : nat-pool-4
Policy                                : nat-policy-4
Started                               : 2016/10/27 13:11:38
Inside router                         : vprn15001
Inside IP address                     : 10.5.15.255

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
*A:PE1#
```

For the 10.5.0.0/27 prefix in VPRN 15002, the subscriber-limit is 64.The 10.5.0.0/27 prefix will be mapped into one outside IP address.

To show the first LSN subscriber for the inside routing instance 15002 of the 10.5.0.0/27 prefix, the following command can be used:

```
*A:PE1# show service nat lsn-subscribers inside-ip 10.5.0.0 inside-router 15002

===============================================================================
NAT LSN subscribers
===============================================================================
Subscriber                  : [LSN-Host@10.5.0.0]
NAT policy                  : nat-policy-4
Subscriber ID               : 276829440
-------------------------------------------------------------------------------
```

```
Type                        : classic-lsn-sub
Inside router               : 15002
Inside IP address prefix    : 10.5.0.0/32
ISA NAT group               : 1
ISA NAT group member        : 1
Outside router              : "Base"
Outside IP address          : 192.168.4.65

-------------------------------------------------------------------------------
No. of LSN subscriber instances: 1
===============================================================================
*A:PE1#
```

To show the LSN blocks corresponding to the first inside IP address within the 10.5.0.0/27 prefix, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.5.0.0 inside-router 15002

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.4.65 [4024..4523]
Pool                                : nat-pool-4
Policy                              : nat-policy-4
Started                             : 2016/10/27 13:12:02
Inside router                       : vprn15002
Inside IP address                   : 10.5.0.0

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
*A:PE1#
```

To show the LSN blocks for the last inside IP address within the 10.5.0.0/27 prefix, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.5.0.31 inside-router 15002

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.4.65 [19524..20023]
Pool                                : nat-pool-4
Policy                              : nat-policy-4
Started                             : 2016/10/27 13:12:02
Inside router                       : vprn15002
Inside IP address                   : 10.5.0.31

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
*A:PE1#
```

Mapping Results

According to this configuration, each inside IP address within VPRN 15001 has one deterministic block of 500 ports and can have up to three dynamic blocks (block-limit =4) of 461 ports each, allowing a maximum of 500+3*461 = 1883 flows.

According to this configuration each inside IP address within VPRN 15002 has one deterministic block of 500 ports and can have up to three dynamic blocks (block-limit =4) of 461 ports each, allowing a maximum of 500+3*461 = 1883 flows.

For VPRN 15002, since the number of LSN subscribers (32) is less than the number of deterministic blocks (64), then 32 deterministic blocks will be wasted, specifically 32*500 = 16,000 ports will be wasted which is not good in terms of capacity planning.

Figure 52 **Case 3 Results**

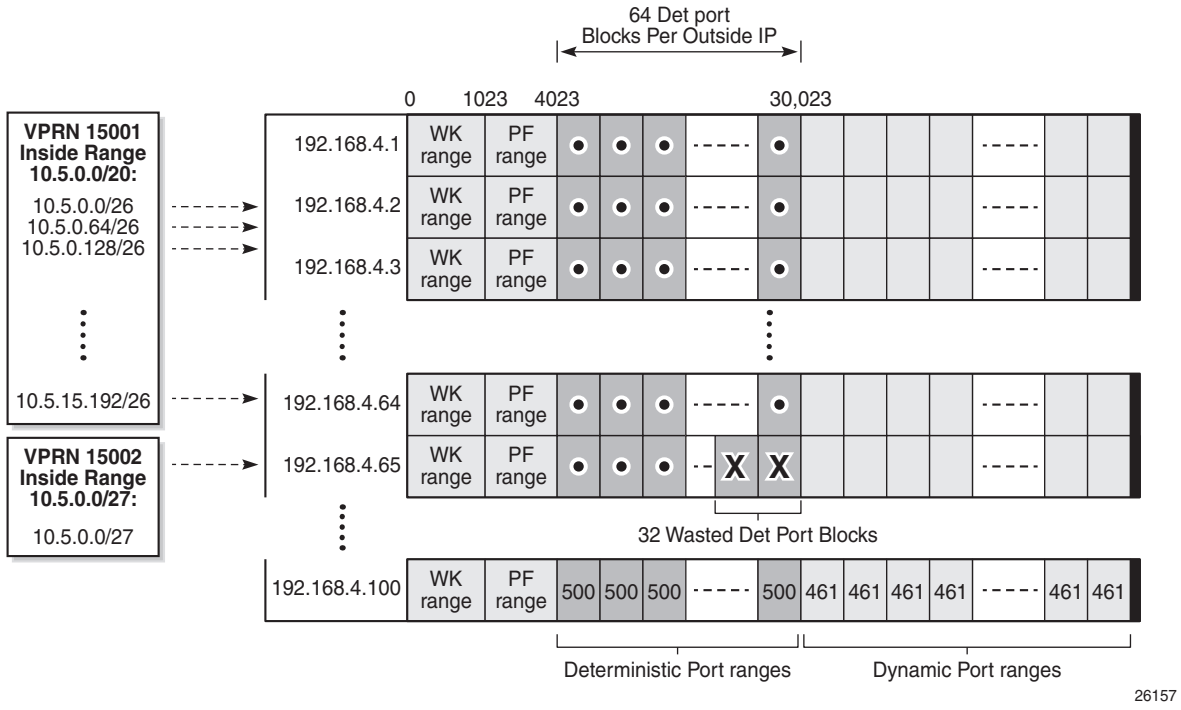

Inverse Mapping

In deterministic LSN44, the inside IP addresses are mapped to outside IP addresses and corresponding port blocks based on a deterministic algorithm. The inverse mapping that reveals the subscriber identity behind the NAT is based on the reversal of this algorithm.

Inverse mappings can be done either online or offline:

- Online — Locally on the 7x50 node, via CLI (MIB)
- Offline — Externally, via a Python script. The purpose of such an offline approach is to provide fast queries without accessing the 7x50.

Figure 53 **Inverse Mapping Approach**

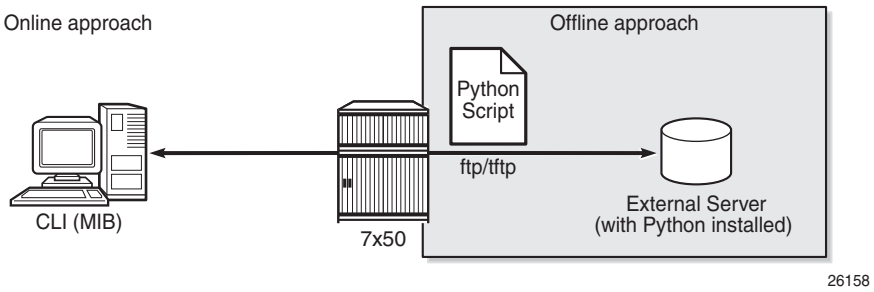

Online Approach

A **tools** command is available which shows the reverse mapping (outside to inside) for deterministic NAT instead of using logging.

```
tools dump nat deterministic-mapping outside-ip <ipv4-address> router <router-
instance> outside-port <[1..65535]>

 <ipv4-address>       : a.b.c.d
 <router-instance>    : <router-name>|<service-id>
                        router-name    - "Base"
                        service-id     - [1..2147483647]
```

Using Case 3 as an example, to obtain (inside IP, inside routing instance), the inverse mapping for a specific (outside IP, outside routing instance, outside port) is done as follows:

```
*A:PE1# tools dump nat deterministic-mapping outside-
ip 192.168.4.1 router "Base" outside-port 4024
classic-lsn-sub inside router 15001 ip 10.5.0.0 --
 outside router Base ip 192.168.4.1 port 4024 at Fri Oct 28 13:04:22 UTC 2016
*A:PE1#


*A:PE1# tools dump nat deterministic-mapping outside-
ip 192.168.4.65 router "Base" outside-port 4024
classic-lsn-sub inside router 15002 ip 10.5.0.0 --
 outside router Base ip 192.168.4.65 port 4024 at Fri Oct 28 13:04:45 UTC 2016
*A:PE1#
```

Offline Approach

The purpose of such an offline approach is to provide fast queries without the need to directly query the 7x50.

This is achieved by generating and exporting a Python script for reverse querying, which is a manual operation that needs to be repeated every time there is configuration change in deterministic NAT.

The script is exported (manually) to the external system.

To configure remotely the location for the Python script, the following command is used:

```
configure service nat deterministic-script location <remote-url>
```

remote-url — A remote location where the script is stored:
[{ftp://|tftp://}<login>:<pswd>@ <remote-locn>/][<file-path>]
Maximum length is 180 characters.

Once the script location is specified, the script can be exported to that location using the following command:

```
admin nat save-deterministic-script
```

Using the following command the status of the script can be checked, and whether it is necessary to re-save (export) the script or not:

```
*A:PE1# show service nat deterministic-script

===============================================================================
Deterministic NAT script data
===============================================================================
Location                    : ftp://*:*@123.123.123.123/pub/python/detnat.py
Save needed                 : no
Last save result            : success
Last save time              : 2016/10/28 13:05:41
===============================================================================
*A:PE1#
```

The external system must have the Python scripting language installed with the following modules: getopt, math, os, socket, and sys.

The Python script can then be run on the external server; the parameters are as follows:

```
[user@123.123.123.123 ~]$ ./detnat.py
Error: need exactly one of --forward or --backward arguments

Usage: detnat.py DIRECTION PARAMETERS
Perform forward or backard NAPT according to the configured deterministic rules.

DIRECTION:
  -f, --forward             Translate from inside to outside address/port
  -b, --backward            Translate from outside to inside address/port

PARAMETERS:
  -a, --address=IP-ADDRESS  The address to translate. IPv6 addresses can be
                            specified in shorthand or full notation.
  -p, --port=PORT           The outside port in case of backward translation.
  -s, --service=SERVICE-ID  The service where the IP-ADDRESS originates from.
                            This is the inside service in case of forward
                            translation and the outside service in case of
                            backward translation.
                            To specify the base router, this option must be
                            omitted.

  -h, --help                Show this help message
[user@123.123.123.123 ~]$
```

where deterministic-nat.py is the name of the python script previously exported.

As an example of a forward query:

```
[user@123.123.123.123 ~]$ ./detnat.py -f -s 15001 -a 10.0.0.1
classic-lsn-
sub has public ip address 192.168.0.1 from base router and is using ports [4324 -
 4623]
[user@123.123.123.123 ~]$
```

As an example of a reverse query:

```
[user@123.123.123.123 ~]$ ./detnat.y -b -s 0 -a 192.168.0.1 -p 4325
classic-lsn-sub has private ip address 10.0.0.1 from service 15001
[user@123.123.123.123 ~]$
```

Simultaneous Support of Deterministic and Non-Deterministic NAT

Deterministic NAT can be used simultaneously with non-deterministic NAT within the same inside routing instance. However, they cannot share the same pool.

An outside pool can be only deterministic (although expandable by dynamic ports blocks) or non-deterministic at any given time (a non-deterministic pool is a pool that contains dynamic port-blocks only).

The following show a configuration using deterministic NAT simultaneously with non-deterministic NAT.

The NAT **outside** pools are configured as follows:

```
configure
    router
        nat
            outside
                pool "nat-pool-1" nat-group 1 type large-scale create
                    port-reservation ports 180
                    port-forwarding-range 4023
                    subscriber-limit 128
                    deterministic
                        port-reservation 300
                    exit
                    address-range 192.168.0.1 192.168.0.100 create
                    exit
                    no shutdown
                exit
                pool "nat-pool-Non-Deterministic" nat-group 1 type large-
scale create
                    address-range 192.168.7.1 192.168.7.100 create
                    exit
                    no shutdown
                exit
            exit
        exit
    exit
exit
```

The NAT **policies** are configured as follows:

```
configure
    service
        nat
            nat-policy "nat-policy-1" create
                block-limit 4
                pool "nat-pool-1" router Base
            exit
            nat-policy "nat-policy-Non-Deterministic" create
                pool "nat-pool-Non-Deterministic" router Base
            exit
        exit
    exit
exit
```

The NAT **inside prefixes** are configured as follows:

```
configure
    service
        vprn 15001 customer 1 create
            nat
                inside
                    destination-prefix 0.0.0.0/0
                    classic-lsn-max-subscriber-limit 256
                    deterministic
                        prefix 10.0.0.0/24 subscriber-type classic-lsn-sub
                                                nat-policy "nat-policy-1" create
                            map start 10.0.0.0 end 10.0.0.255 to 192.168.0.1
                            no shutdown
                        exit
                    exit
                    nat-policy "nat-policy-Non-Deterministic"
                exit
            exit
            no shutdown
        exit
    exit
exit
```

In this example, the inside IP prefixes that do not match any of the deterministic prefixes will be NATed using a non-deterministic pool.

Figure 54 **Sending Flows: Deterministic + non-Deterministic NAT**

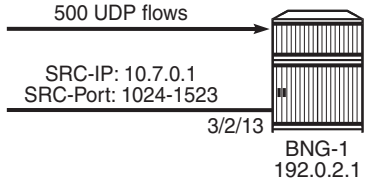

To check which NAT pool/NAT policy is used for NATing the inside IP 10.7.0.1, the following command can be used:

```
*A:PE1# show router nat lsn-blocks inside-ip 10.7.0.1

===============================================================================
Large-Scale NAT blocks for Base
===============================================================================
192.168.7.100 [1024..1527]
Pool                                : nat-pool-Non-Deterministic
Policy                              : nat-policy-Non-Deterministic
Started                             : 2016/10/28 13:24:56
Inside router                       : vprn15001
Inside IP address                   : 10.7.0.1

-------------------------------------------------------------------------------
Number of blocks: 1
===============================================================================
*A:PE1#
```

Conclusion

This example provides the commands required for configuring deterministic LSN44 NAT. Both deterministic as well as non-deterministic NAT are supported, with simultaneous operation being possible.

Inverse query can be done online or offline to retrieve the NAT mappings. Logging is not needed as long as there are no dynamic blocks assigned to LSN44 subscribers.

IP/GRE Termination

This chapter describes advanced IP/GRE termination configurations.

Topics in this chapter include:

Applicability

The chapter was initially written for SR OS release 9.0.R8. The CLI in the current edition corresponds to 14.0.R5.

The IP GRE tunnel termination configuration described in this chapter requires an MS-ISA equipped on IOM2-20g or IOM3-XP. IP GRE tunnels without ISA are beyond the scope of this chapter.

Chassis mode B or higher is required.

Summary

Initially, the 7750 SR only supported GRE SDP tunnels which use pseudowire encapsulation. From SR OS 8.0.R5 onward, the 7750 SR supports tunneling IPv4 packets in an IPv4 GRE tunnel. A common use case is remote access to a VPRN over a public IP network because IP/GRE tunneling allows encapsulated packets to follow a path based on the outer IP header which is useful when the inner IP packet cannot or should not be forwarded natively over this path.

This chapter provides configuration and troubleshooting commands for IP/GRE termination.

Overview

Generic Routing Encapsulation (GRE) allows packets of one protocol, the payload protocol, to be encapsulated by packets of another protocol, called the delivery protocol. A GRE packet has an Outer Delivery Header, GRE Header and Payload Packet (Figure 55).

Figure 55 **GRE Packet Format**

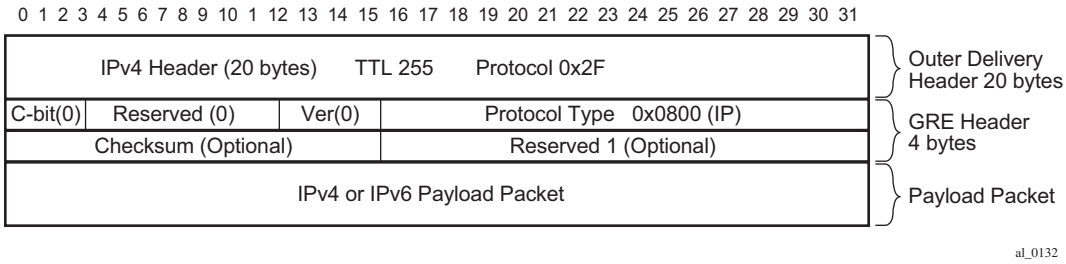

The outer delivery and GRE header for outgoing traffic is as follows.

- Outer Delivery header
  - The source address in the IPv4 delivery header is the configured source address.
  - The destination address in the IPv4 delivery header is the configured remote-ip (or backup-remote-ip) address.
  - The IP protocol value in the IPv4 delivery header is 0x2F or 47 (GRE).
  - The DSCP in the IPv4 Outer Delivery header is:
    - Set to the value configured for the tunnel.
    - Otherwise, the DSCP value from the Payload Packet is copied into the Outer Delivery header.
  - The TTL in the IPv4 Outer Delivery header is set to 255.
- GRE Header
  - The Checksum (C) bit in the GRE header is set to 0 (no checksum present).
  - The version in the GRE header is 0.
  - The protocol type in the GRE header is 0x0800 for IPv4.

The outer delivery and GRE header for incoming traffic is as follows:

- Outer Delivery header
  - If the packet is a fragment (More Fragments=1, non-zero fragment offset), it is dropped.

- If the Checksum (C) bit in the GRE header is set, then the included checksum is validated; if the checksum is incorrect, the packet is discarded.
- If the version in the GRE header is not 0, the packet is discarded.
- If the source/destination address pair in the IPv4 delivery header is any other combination, the packet is dropped.

• GRE Header

- If the Checksum (C) bit in the GRE header is set, then the included checksum is validated; if the checksum is incorrect, the packet is discarded.
- If the version in the GRE header is not 0, the packet is discarded.

7750 SR Implementation

Encapsulation, de-encapsulation and other datapath operations related to IP/GRE are handled by the isa-tunnel MDA.

For GRE tunnels configured as SDPs (which are not covered by this section), no isa-tunnel MDA is required.

Figure 56 **7x50 Implementation**

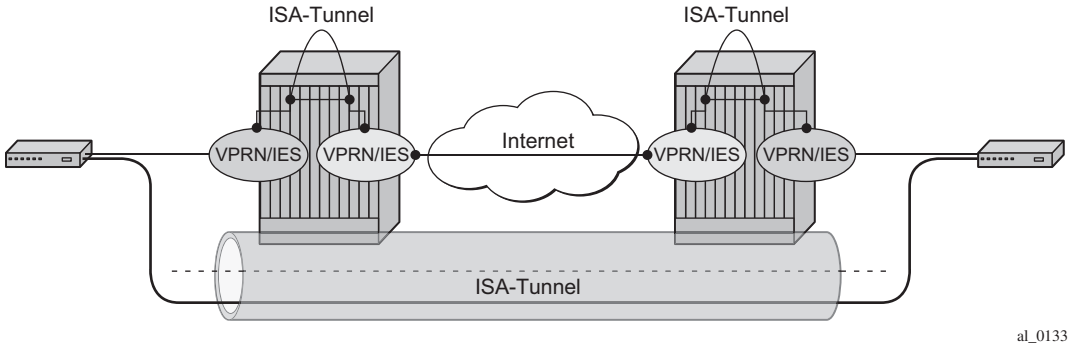

The 7750 SR initially supported the IP/GRE tunnels with static routes and BGP. IP/GRE tunnels have been enhanced by adding support for OSPF and BFD on private tunnel interfaces (used with static routes, OSPF or BGP) and GRE protection by tunneling into an IPSec tunnel.

Figure 57 **IP/GRE over IPSec Tunnel**

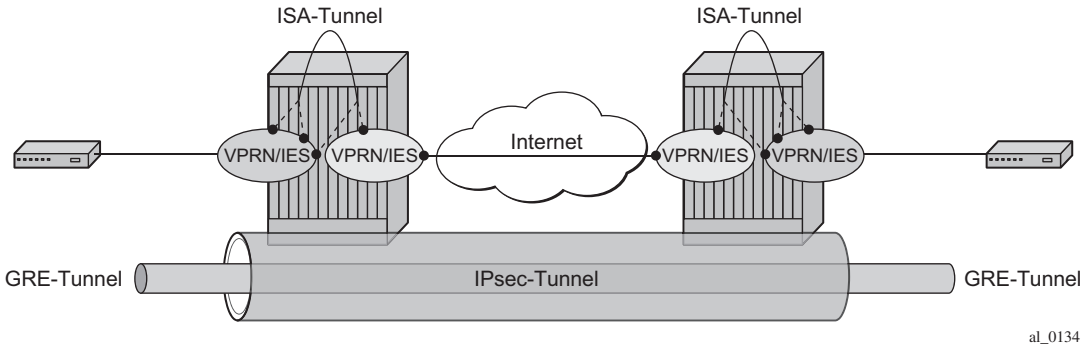

Configuration

Tunnel ISA MDA

The isa-tunnel MDA supports IP/GRE and IPSec tunnels.

```
*A:PE-1#configure card 1 mda 2 mda-type "isa-tunnel"
```

To check the MDA configuration:

```
*A:PE-1# show mda

===============================================================================
MDA Summary
===============================================================================
Slot  Mda   Provisioned Type                              Admin     Operational
               Equipped Type (if different)              State     State
-------------------------------------------------------------------------------
1     1     m4-10gb-xp-xfp                               up        up
      2     isa-tunnel                                   up        up
                isa-ms
```

Tunnel Groups and Tunnel Group Restrictions

The first step of the GRE tunnel configuration is to configure a tunnel-group.

A tunnel group can have one tunnel ISA designated primary and optionally one tunnel-ISA designated backup. When a GRE tunnel is created, it is assigned to the primary tunnel-ISA in its tunnel group. If the primary tunnel-ISA fails, the backup tunnel-ISA (if not already claimed by another tunnel-group) takes over for the failed card.

```
*A:PE-1# configure isa tunnel-group
  - tunnel-group <tunnel-group-id> [create]
  - tunnel-group <tunnel-group-id> isa-scale-mode <isa-scale-mode> [create]
  - no tunnel-group <tunnel-group-id>

 <tunnel-group-id>    : [1..16]
 <isa-scale-mode>     : tunnel-limit-2k, k = 1024
 <create>             : keyword - mandatory while creating an entry.

 [no] active-mda-num* - Configure number of active MDAs
 [no] backup          - Configure ISA-Tunnel-Group backup ISA
 [no] description     - Configure the ISA group description
 [no] ipsec-responde* - Enable/Disable responder-only for IPsec Ikev2 tunnels only
 [no] mda             - Configure MDA to associate with
 [no] multi-active    - Configure multi-active status of tunnel-group
 [no] primary         - Configure ISA-Tunnel-Group primary ISA
 [no] reassembly      - Configure reassembly wait time
 [no] shutdown        - Administratively enable/disable an ISA-Tunnel-Group

*A:PE-1# configure
    isa
        tunnel-group 1 create
            primary 1/2
            backup 2/1
            no shutdown
        exit
```

The failed tunnels are re-established using a cold-standby on the backup tunnel-ISA (cold standby means the backup tunnel-ISA has no state or configuration information about the tunnels prior to the failure).

A tunnel-ISA cannot be primary for more than one tunnel group.

```
*A:PE-1>config>isa# tunnel-group 2 create
*A:PE-1>config>isa>tunnel-grp# primary 1/2
MINOR: IPSECGRPMGR #1003 The specified MDA is primary in another Tunnel Group
```

A tunnel-ISA cannot be primary in one tunnel group and backup in another tunnel group.

```
*A:PE-1>config>isa# tunnel-group 2 create
*A:PE-1>config>isa>tunnel-grp# backup 1/2
MINOR: IPSECGRPMGR #1003 The specified MDA is primary in another Tunnel Group
```

To check the ISA tunnel-group (after tunnel group 2 has been removed):

```
*A:PE-1# show isa tunnel-group
===============================================================================
```

```
ISA Tunnel Groups
===============================================================================
Tunnel    PrimaryIsa             BackupIsa    ActiveIsa    Admin     Oper
GroupId                                                    State     State
-------------------------------------------------------------------------------
1         1/2                    2/1          1/2          Up        Up
-------------------------------------------------------------------------------
No. of ISA Tunnel Groups: 1
===============================================================================
*A:PE-1#
```

To check the number of the IP (GRE) tunnels (after configuring IES and VPRN services with tunnel interfaces):

```
*A:PE-1# show ip tunnel count
-------------------------------------------------------------------------------
IP Tunnels: 2
-------------------------------------------------------------------------------
*A:PE-1#
```

To check all IP tunnels:

```
*A:PE-1# show ip tunnel

===============================================================================
IP Tunnels
===============================================================================
TunnelName                     SapId                        SvcId      Admn
 Local Address                                              DlvrySvcId Oper
  OperRemoteAddress
-------------------------------------------------------------------------------
gre-tunnel-1                   tunnel-1.private:1           1          Up
 192.168.1.1                                                2          Up
  192.168.3.1
protected-gre-tunnel           tunnel-1.private:5           3          Up
 192.168.11.1                                               3          Up
  192.168.33.1
-------------------------------------------------------------------------------
IP Tunnels: 2
===============================================================================
*A:PE-1#
```

To check the detailed tunnel information:

```
*A:PE-1# show ip tunnel "gre-tunnel-1"

===============================================================================
IP Tunnel Configuration Detail
===============================================================================
Service Id       : 1                    Sap Id           : tunnel-1.private:1
Tunnel Name      : gre-tunnel-1
Description      : None
GRE Header       : Yes                  Delivery Service : 2
GRE Keys Set     : False
GRE Send Key     : N/A                  GRE Receive Key  : N/A
Admin State      : Up                   Oper State       : Up
```

```
Source Address   : 192.168.1.1
Remote Address   : 192.168.3.1
Backup Address   : (Not Specified)
Oper Remote Addr : 192.168.3.1
DSCP             : None
Reassembly       : inherit
Clear DF Bit     : false                IP MTU           : max
Encap IP MTU     : max
Pkt Too Big      : true
Pkt Too Big Numb*: 100                  Pkt Too Big Intvl: 10 secs
Oper Flags       : None
Last Oper Changed: 10/26/2016 08:43:56
Host MDA         : 1/2

-------------------------------------------------------------------------------
Target Address Table
-------------------------------------------------------------------------------
Destination IP                          IP Resolved Status
-------------------------------------------------------------------------------
10.0.0.2                                Yes
-------------------------------------------------------------------------------

===============================================================================
IP Tunnel Statistics: gre-tunnel-1
===============================================================================
Errors Rx        : 0                    Errors Tx        : 0
Pkts Rx          : 145                  Pkts Tx          : 145
Bytes Rx         : 9394                 Bytes Tx         : 9599
Key Ignored Rx   : 0                    Too Big Tx       : 0
Seq Ignored Rx   : 0
Vers Unsup. Rx   : 0
Invalid Chksum Rx: 0
Key Mismatch Rx  : 0
===============================================================================

===============================================================================
Fragmentation Statistics
===============================================================================
Encapsulation Overhead                : 24
Pre-Encapsulation
    Fragmentation Count               : 0
    Last Fragmented Packet Size       : 0
Post-Encapsulation
    Fragmentation Count               : 0
    Last Fragmented Packet Size       : 0
===============================================================================
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:PE-1#
```

Interfaces

The interface toward the Internet (or WAN):

- Can be a network interface or VPRN/IES interface.
- Provides IP reachability.

The tunnel public interface:

- Can be an IES or VPRN interface.
- Represents the public side of the tunnel-ISA.

The delivery VPRN/IES service (the service connected to the Internet) must have at least one IP interface associated with a public tunnel SAP in order to receive and process GRE encapsulated packets.

The public tunnel SAP type has the format **tunnel**-*id*.**private** | **public**:*tag* (where the *id* corresponds to the tunnel group) as shown in the following example.

```
*A:PE-1>config>service>ies>if# sap ?
 ---snip---
 <sap-id>
---snip---
                        tunnel-id      - tunnel-<id>.<private|public>:<tag>
                          tunnel         - keyword
                          id             - [1..16]
                          tag            - [0..4094]

*A:PE-1# configure
    service
        ies 2 customer 1 create
            interface "int-tunnel-public" create
                address 192.168.1.2/30
                tos-marking-state untrusted
                sap tunnel-1.public:1 create
                exit
            exit
            interface "int-PE-1-CE-3" create
                address 192.168.13.1/24
                sap 1/1/2:2 create
                exit
            exit
            no shutdown
        exit
```

PE-1 has address 192.168.1.2/30 assigned to the interface “int-tunnel-public in IES 2. In a similar way, CE-3 has address 192.168.3.2/30 assigned to the interface “int-tunnel-public” in IES 2.

In order to reach the remote IP address from PE-1, a static route is configured, as follows:

```
*A:PE-1# configure router static-route-entry 192.168.3.0/30 next-
hop 192.168.13.2 no shutdown
```

In a similar way, a static route is configured on CE-3 to reach 192.168.1.0/30.

Mask /32 is not supported on the public tunnel. When address 192.168.1.2/32 is configured on the interface, the public tunnel cannot be created, as follows:

```
*A:PE-1>config>service>ies>if# address 192.168.1.2/32
*A:PE-1>config>service>ies# interface "tunnel-public" sap tunnel-1.public:1 create
INFO: PIP #1288 Cannot bind when there are /32 or /128 addresses configured
```

Therefore, the address configured on the interface will have mask /30 instead of /32, as shown earlier.

The tunnel private interface:

- Can be an IES or VPRN interface.
- Represents the private side of the tunnel-ISA.

The private tunnel SAP has the format **tunnel**-*id*.**private** | **public**:*tag* (where the *id* corresponds to the tunnel-group) as shown in the following example where an unprotected GRE tunnel is configured under the SAP.

```
*A:PE-1>config>service>vprn>if# sap ?
---snip---
 <sap-id>
---snip---
                        tunnel-id      - tunnel-<id>.<private|public>:<tag>
                          tunnel         - keyword
                          id             - [1..16]
                          tag            - [0..4094]

*A:PE-1# configure
    service
        vprn 1 customer 1 create
            route-distinguisher 64496:1
            vrf-target target:64496:1
            interface "int-gre-tunnel" tunnel create
                address 10.0.0.1/30
                sap tunnel-1.private:1 create
                    ip-tunnel "gre-tunnel-1" create
                        gre-header
                        dest-ip 10.0.0.2
                        ---snip---
```

It is not mandatory to have the same tag (internal dot1q) in private and public GRE tunnels.

```
    sap tunnel-1.private:1 <=> sap tunnel-1.public:2
```

Unprotected GRE Tunnel Configuration

To associate an unprotected GRE tunnel with a private tunnel SAP, the **ip-tunnel** command is configured in the SAP context.

```
*A:PE-1# configure
    service
        vprn 1 customer 1 create
        ---snip---
            interface "int-gre-tunnel" tunnel create
                address 10.0.0.1/30
                sap tunnel-1.private:1 create
                    ip-tunnel "gre-tunnel-1" create
                        dest-ip 10.0.0.2
                        gre-header
                        ---snip---
```

The **dest-ip** keyword followed by the private IP address of the remote tunnel endpoint is mandatory.

If this remote IP address is not within the subnet of the local private endpoint, then the tunnel will not come up.

Under the **ip-tunnel** context, configure the following parameters:

- The source address of the GRE tunnel. This is the source IPv4 address of GRE encapsulated packets sent by the delivery service. It must be an address in the subnet of the associated public tunnel SAP interface.
- The remote IP address. If this address is reachable in the delivery service (there is a route), then this is the destination IPv4 address of GRE encapsulated packets sent by the delivery service.
- The backup remote IP address. If the remote IP address of the tunnel is not reachable, then this is the destination IPv4 address of GRE encapsulated packets sent by the delivery service.
- The delivery service. This is the identifier or name of the IES or VPRN service where GRE encapsulated packets are injected and terminated. The delivery service can be the same service where the private tunnel SAP interface resides.
- The DSCP marking in the outer IP header of GRE encapsulated packets. If this is not configured, then the default copies the DSCP from the inner IP header to the outer IP header.

```
*A:PE-1# configure service
        vprn 1 customer 1 create
            route-distinguisher 64496:1
            vrf-target target:64496:1
            interface "int-gre-tunnel" tunnel create
                address 10.0.0.1/30
                sap tunnel-1.private:1 create
```

```
ip-tunnel "gre-tunnel-1" create
    dest-ip 10.0.0.2
    gre-header
    source 192.168.1.1
    remote-ip 192.168.3.1
    delivery-service 2
    dscp af22
    no shutdown
exit
---snip---
```

- A private tunnel SAP can have only one IP/GRE tunnel (per SAP).

```
*A:PE-1# configure service vprn 1 interface "int-gre-tunnel" sap tunnel-
1.private:1 ip-tunnel "gre-tunnel-2" create
MINOR: SVCMGR #5120 Only one IP tunnel allowed per SAP
```

IP/GRE Tunneling via Static Route

A static route can reference the GRE tunnel directly (by next-hop IP address) or the GRE tunnel can be the resolved next-hop for an indirect static route (Figure 58).

Figure 58 **GRE for Remote Access to a VPRN Service**

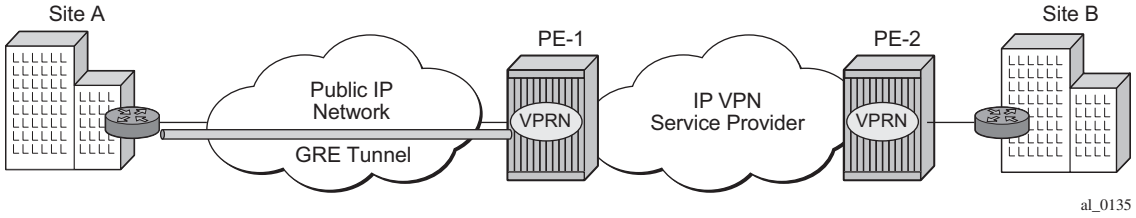

The details of both ends on the GRE tunnel, at site A and PE-1, are shown in Figure 59. The node at left hand side is CE-3 at site A.

Figure 59 **IP/GRE Tunneling via Static Route**

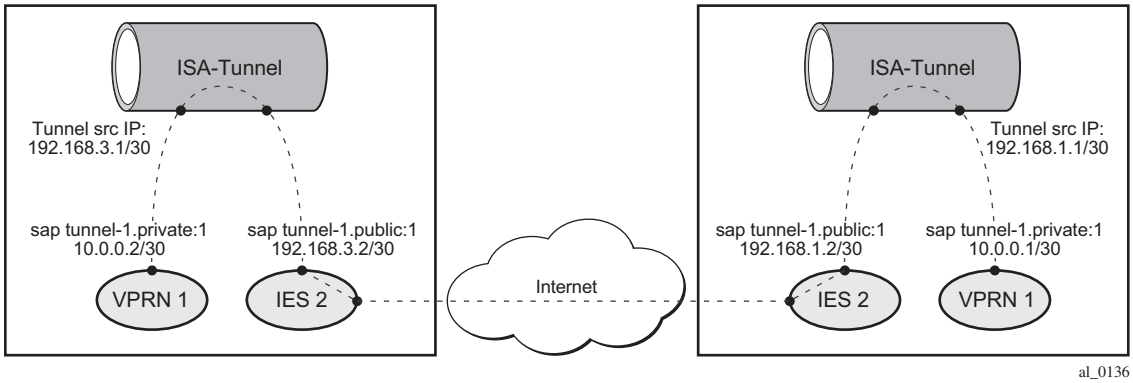

The following shows the configuration of VPRN 1 on PE-1.

```
*A:PE-1# configure
    service
        vprn 1 customer 1 create
            route-distinguisher 64496:1
            vrf-target target:64496:1
            interface "int-gre-tunnel" tunnel create
                address 10.0.0.1/30
                sap tunnel-1.private:1 create
                    ip-tunnel "gre-tunnel-1" create
                        dest-ip 10.0.0.2
                        gre-header
                        source 192.168.1.1
                        remote-ip 192.168.3.1
                        delivery-service 2
                        no shutdown
                    exit
                exit
            exit
            static-route-entry 172.16.3.1/32
                next-hop 10.0.0.2
                    no shutdown
                exit
            exit
            no shutdown
        exit
```

The configuration of the VPRN on CE-3 is similar.

To check the static route status:

```
*A:PE-1# show router 1 static-route

===============================================================================
Static Route Table (Service: 1)  Family: IPv4
===============================================================================
Prefix                                      Tag         Met    Pref Type Act
   Next Hop                                  Interface
-------------------------------------------------------------------------------
```

```
172.16.3.1/32                                  0           1      5    NH   Y
   10.0.0.2                                     int-gre-tunnel
------------------------------------------------------------------------------
No. of Static Routes: 1
==============================================================================
*A:PE-1#
```

IP/GRE Tunneling via BGP Peering

In this section, the configuration has BGP running inside the GRE tunnel.

```
*A:PE-1# configure
    service
        vprn 1 customer 1 create
            autonomous-system 64496
            route-distinguisher 64496:1
            vrf-target target:64496:1
            interface "int-gre-tunnel" tunnel create
                address 10.0.0.1/30
                sap tunnel-1.private:1 create
                    ip-tunnel "gre-tunnel-1" create
                        dest-ip 10.0.0.2
                        gre-header
                        source 192.168.1.1
                        remote-ip 192.168.3.1
                        delivery-service 2
                        no shutdown
                    exit
                exit
            exit
            interface "loopback1" create
                address 172.31.1.1/32
                loopback
            exit
            static-route-entry 172.16.3.1/32
                next-hop 10.0.0.2
                    no shutdown
                exit
            exit
            bgp
                group "group-1"
                    type internal
                    local-address 172.31.1.1
                    neighbor 172.16.3.1
                    exit
                exit
                no shutdown
            exit
            no shutdown
        exit
```

It is mandatory to configure the autonomous-system under the VPRN context, otherwise the BGP neighboring will not be established.

The configuration of the VPRN on CE-3 is similar.

To check the BGP status:

```
*A:PE-1# show router 1 bgp neighbor

===============================================================================
BGP Neighbor
===============================================================================
-------------------------------------------------------------------------------
Peer                 : 172.16.3.1
Description          : (Not Specified)
Group                : group-1
-------------------------------------------------------------------------------
Peer AS              : 64496            Peer Port            : 179
Peer Address         : 172.16.3.1
Local AS             : 64496            Local Port           : 49527
Local Address        : 172.31.1.1
Peer Type            : Internal         Dynamic Peer         : No
State                : Established      Last State           : Active
---snip---
```

IP/GRE Tunneling via OSPFv2 Peering

OSPF can be run on IES and VPRN IP interfaces associated with private IP/GRE tunnel SAPs.

All OSPF features are supported, including area 0 and non-area 0 support, virtual links, authentication, BFD, configurable protocol timers.

```
*A:PE-1# configure
     service
        vprn 1 customer 1 create
            route-distinguisher 64496:1
            vrf-target target:64496:1
            interface "int-gre-tunnel" tunnel create
                address 10.0.0.1/30
                sap tunnel-1.private:1 create
                    ip-tunnel "gre-tunnel-1" create
                        dest-ip 10.0.0.2
                        gre-header
                        source 192.168.1.1
                        remote-ip 192.168.3.1
                        delivery-service 2
                        no shutdown
                    exit
                exit
            exit
            interface "loopback1" create
                address 172.31.1.1/32
                loopback
            exit
            ospf
```

```
                area 0.0.0.0
                    interface "int-gre-tunnel"
                    exit
                    interface "loopback1"
                    exit
                exit
                no shutdown
            exit
            no shutdown
        exit
```

The following command shows the OSPF neighbors for VPRN 1:

```
*A:PE-1# show router 1 ospf neighbor

===============================================================================
Rtr vprn1 OSPFv2 Instance 0 Neighbors
===============================================================================
Interface-Name                   Rtr Id          State      Pri  RetxQ   TTL
   Area-Id
-------------------------------------------------------------------------------
int-gre-tunnel                   192.0.2.3       Full       1    0       33
   0.0.0.0
-------------------------------------------------------------------------------
No. of Neighbors: 1
===============================================================================
*A:PE-1#
```

The OSPF routes in the routing table of VPRN 1 are as follows:

```
*A:PE-1# show router 1 route-table protocol ospf

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
172.16.3.1/32                                   Remote  OSPF      00h05m46s  10
       10.0.0.2                                                        10
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
*A:PE-1#
```

IP/GRE Tunneling Protection using IPSec Tunnel Mode

To provide protection against the potential threats (such as spoofing) the GRE packets can be encrypted and authenticated using IPSec.

GRE packets receive IPSec protection by forwarding them, after encapsulation by a tunnel-ISA, into an IPSec tunnel supported by another (or the same) tunnel-ISA.

Note that when configuring GRE protection by an IPSec tunnel:

- A GRE tunnel and its protecting IPSec tunnel may belong to the same or different tunnel-groups (the same tunnel-group is assumed in the example below).
- A GRE tunnel and its protecting IPSec tunnel may be assigned to the same tunnel-ISA (if they belong to the same tunnel-group) or different tunnel-ISAs.
- A single IPSec tunnel can protect one or more GRE tunnels in addition to other IP traffic that meets the IPSec security policy.
- The private IPSec tunnel SAP interface and public GRE tunnel SAP interface are always part of the same VPRN. The private GRE tunnel SAP interface can be part of this same VPRN or a different VPRN.

In the following example, the GRE tunnel and its protecting IPSec tunnel belong to the same tunnel group.

Figure 60 **Example GRE over IPSec Tunnel**

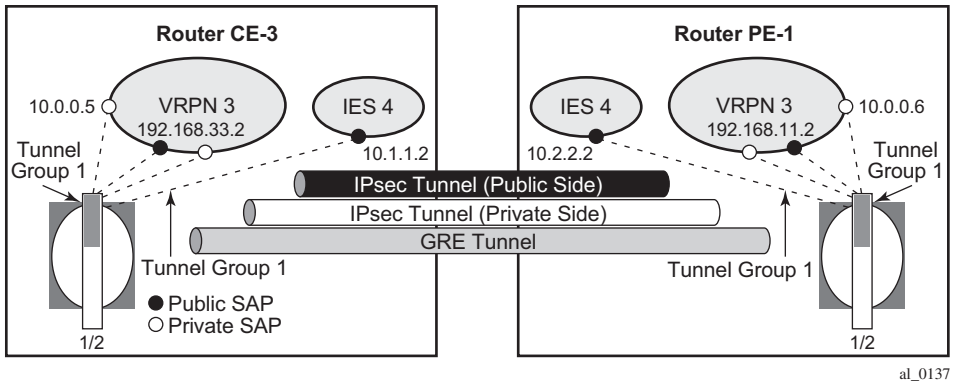

IPSec Configuration

An ike-policy and ipsec-transform must be configured on PE-1 and CE-3, as follows:

```
configure
    ipsec
        ike-policy 1 create
            dh-group 5
        exit
        ipsec-transform 1 create
            esp-encryption-algorithm aes256
        exit
```

The public/private side of the GRE tunnel and the private side of the IPSec tunnel are in the same VPRN, as shown in the following configuration example.

```
*A:PE-1# configure
    service
        vprn 3 customer 1 create
            ipsec
                security-policy 1 create
                    entry 1 create
                        local-ip 192.168.11.0/24
                        remote-ip 192.168.33.0/24
                    exit
                exit
            exit
            route-distinguisher 64496:3
            vrf-target target:64496:3
            interface "int-private-ipsec-1" tunnel create
                sap tunnel-1.private:3 create
                    ipsec-tunnel "ipsec-tunnel-for-gre-tunnel" create
                        security-policy 1
                        local-gateway-address 10.2.2.1 peer 10.1.1.1
                                                        delivery-service 4
                        dynamic-keying
                            ike-policy 1
                            pre-shared-key "pass"
                            transform 1
                        exit
                        no shutdown
                    exit
                exit
            exit
            interface "int-public-gre-1" create
                address 192.168.11.2/24
                sap tunnel-1.public:4 create
                exit
            exit
            interface "int-private-gre-1" tunnel create
                address 10.0.0.6/30
                sap tunnel-1.private:5 create
                    ip-tunnel "protected-gre-tunnel" create
                        dest-ip 10.0.0.5
                        gre-header
                        source 192.168.11.1
```

```
                    remote-ip 192.168.33.1
                    delivery-service 3
                    no shutdown
                exit
            exit
        exit
        static-route-entry 192.168.33.0/24
            ipsec-tunnel "ipsec-tunnel-for-gre-tunnel"
                no shutdown
            exit
        exit
        no shutdown
    exit
```

The following displays a configuration example of the public side of the IPSec tunnel:

```
*A:PE-1# configure
    service
        ies 4 customer 1 create
            interface "public-ipsec-1" create
                address 10.2.2.2/24
                tos-marking-state untrusted
                sap tunnel-1.public:3 create
                exit
            exit
            interface "int2-PE-1-CE-3" create
                address 192.168.113.1/30
                sap 1/1/2:4 create
                exit
            exit
            no shutdown
        exit
```

The following static route is configured in the base router on PE-1:

```
*A:PE-1# configure router static-route-entry 10.1.1.0/24 next-
hop 192.168.113.2 no shutdown
```

The configuration is similar on CE-3.

The following command shows that the tunnel “protected-gre-tunnel” with SAP tunnel-1.private:5 is up:

```
*A:PE-1# show ip tunnel

===============================================================================
IP Tunnels
===============================================================================
TunnelName                       SapId                          SvcId      Admn
 Local Address                                                  DlvrySvcId Oper
  OperRemoteAddress
-------------------------------------------------------------------------------
gre-tunnel-1                     tunnel-1.private:1             1          Up
 192.168.1.1                                                    2          Up
  192.168.3.1
```

```
protected-gre-tunnel             tunnel-1.private:5             3          Up
 192.168.11.1                                                  3          Up
  192.168.33.1
-------------------------------------------------------------------------------
IP Tunnels: 2
===============================================================================
*A:PE-1#
```

The following command shows the IP/GRE tunnel information for this IPSec tunnel:

```
*A:PE-1# show ip tunnel "protected-gre-tunnel"
===============================================================================
IP Tunnel Configuration Detail
===============================================================================
Service Id       : 3                    Sap Id           : tunnel-1.private:5
Tunnel Name      : protected-gre-tunnel
Description      : None
GRE Header       : Yes                  Delivery Service : 3
GRE Keys Set     : False
GRE Send Key     : N/A                  GRE Receive Key  : N/A
Admin State      : Up                   Oper State       : Up
Source Address   : 192.168.11.1
Remote Address   : 192.168.33.1
Backup Address   : (Not Specified)
Oper Remote Addr : 192.168.33.1
DSCP             : None
Reassembly       : inherit
Clear DF Bit     : false                IP MTU           : max
Encap IP MTU     : max
Pkt Too Big      : true
Pkt Too Big Numb*: 100                  Pkt Too Big Intvl: 10 secs
Oper Flags       : None
Last Oper Changed: 10/26/2016 10:20:08
Host MDA         : 1/2
-------------------------------------------------------------------------------
Target Address Table
-------------------------------------------------------------------------------
Destination IP                          IP Resolved Status
-------------------------------------------------------------------------------
10.0.0.5                                Yes
-------------------------------------------------------------------------------
===============================================================================
IP Tunnel Statistics: protected-gre-tunnel
===============================================================================
Errors Rx        : 0                    Errors Tx        : 0
Pkts Rx          : 0                    Pkts Tx          : 0
Bytes Rx         : 0                    Bytes Tx         : 0
Key Ignored Rx   : 0                    Too Big Tx       : 0
Seq Ignored Rx   : 0
Vers Unsup. Rx   : 0
Invalid Chksum Rx: 0
Key Mismatch Rx  : 0
===============================================================================
===============================================================================
Fragmentation Statistics
===============================================================================
Encapsulation Overhead                : 24
```

```
Pre-Encapsulation
    Fragmentation Count                : 0
    Last Fragmented Packet Size        : 0
Post-Encapsulation
    Fragmentation Count                : 0
    Last Fragmented Packet Size        : 0
===============================================================================
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:PE-1#
```

By default, the IPSec tunnel is down if it is not used by any traffic, as follows:

```
*A:PE-1# show ipsec tunnel
===============================================================================
IPsec Tunnels
===============================================================================
TunnelName                        LocalAddress      SvcId        Admn   Keying
  SapId                             RemoteAddress     DlvrySvcId   Oper   Sec
                                                                          Plcy
-------------------------------------------------------------------------------
ipsec-tunnel-for-gre-tunnel       10.2.2.1          3            Up     Dynamic
  tunnel-1.private:3                10.1.1.1          4              Down   1
-------------------------------------------------------------------------------
```

After it is used by traffic, the status will be changed to be up.

```
*A:PE-1# ping router 3 10.0.0.5
PING 10.0.0.5 56 data bytes
64 bytes from 10.0.0.5: icmp_seq=1 ttl=64 time=1.42ms.
64 bytes from 10.0.0.5: icmp_seq=2 ttl=64 time=1.35ms.
64 bytes from 10.0.0.5: icmp_seq=3 ttl=64 time=1.26ms.
64 bytes from 10.0.0.5: icmp_seq=4 ttl=64 time=1.34ms.
64 bytes from 10.0.0.5: icmp_seq=5 ttl=64 time=1.28ms.

---- 10.0.0.5 PING Statistics ----
5 packets transmitted, 5 packets received, 0.00% packet loss
round-trip min = 1.26ms, avg = 1.33ms, max = 1.42ms, stddev = 0.058ms
*A:PE-1#
```

The IPSec tunnel is now up, as follows:

```
*A:PE-1# show ipsec tunnel
===============================================================================
IPsec Tunnels
===============================================================================
TunnelName                        LocalAddress      SvcId        Admn   Keying
  SapId                             RemoteAddress     DlvrySvcId   Oper   Sec
                                                                          Plcy
-------------------------------------------------------------------------------
ipsec-tunnel-for-gre-tunnel       10.2.2.1          3            Up     Dynamic
  tunnel-1.private:3                10.1.1.1          4              Up     1
-------------------------------------------------------------------------------
IPsec Tunnels: 1
===============================================================================
```

BFD Support on Private Tunnel Interfaces

BFD is supported on IP interfaces associated with private IP/GRE tunnel SAPs. The BFD state of the interface can be used by static routes, OSPFv2 and/or BGP configured on the interface. It is not used to trigger a switchover to the backup remote IP address of the GRE tunnel.

The following displays a static route example:

```
*A:PE-1# configure
    service
        vprn 1 customer 1 create
            route-distinguisher 64496:1
            vrf-target target:64496:1
            interface "int-gre-tunnel" tunnel create
                address 10.0.0.1/30
                bfd 100 receive 100 multiplier 3
                sap tunnel-1.private:1 create
                    ip-tunnel "gre-tunnel-1" create
                        dest-ip 10.0.0.2
                        gre-header
                        source 192.168.1.1
                        remote-ip 192.168.3.1
                        delivery-service 2
                        no shutdown
                    exit
                exit
            exit
            static-route-entry 172.16.3.1/32
                next-hop 10.0.0.2
                    bfd-enable
                    no shutdown
                exit
            exit
            no shutdown
        exit
```

The following command shows that the BFD session on interface “int-gre-tunnel” is up for protocol static:

```
*A:PE-1# show router 1 bfd session

===============================================================================
Legend:  wp = Working path   pp = Protecting path
===============================================================================
BFD Session
===============================================================================
If/Lsp Name/Svc-Id/RSVP-sess  State                Tx Intvl  Rx Intvl  Multipl
  Rem Addr/Info/SdpId:VcId    Protocols            Tx Pkts   Rx Pkts   Type
   LAG port                      LAG ID
-------------------------------------------------------------------------------
int-gre-tunnel                Up                   100       100       3
   10.0.0.2                   static               2976      2954      central
-------------------------------------------------------------------------------
No. of BFD sessions: 1
```

```
===============================================================================
*A:PE-1#
```

When no static routes are configured and OSPF is configured instead, the configuration of VPRN 1 on PE-1 is as follows:

```
*A:PE-1# configure
    service
        vprn 1 customer 1 create
            route-distinguisher 64496:1
            vrf-target target:64496:1
            interface "int-gre-tunnel" tunnel create
                address 10.0.0.1/30
                bfd 100 receive 100 multiplier 3
                sap tunnel-1.private:1 create
                    ip-tunnel "gre-tunnel-1" create
                        dest-ip 10.0.0.2
                        gre-header
                        source 192.168.1.1
                        remote-ip 192.168.3.1
                        delivery-service 2
                        no shutdown
                    exit
                exit
            exit
            interface "loopback1" create
                address 172.31.1.1/32
                bfd 100 receive 100 multiplier 3
                loopback
            exit
            ospf
                area 0.0.0.0
                    interface "int-gre-tunnel"
                        bfd-enable
                        no shutdown
                    exit
                    interface "loopback1"
                        no shutdown
                    exit
                exit
                no shutdown
            exit
            no shutdown
```

The following shows that the BFD session is up for protocol OSPF on interface “int-gre-tunnel”:

```
*A:PE-1# show router 1 bfd session

===============================================================================
Legend:  wp = Working path   pp = Protecting path
===============================================================================
BFD Session
===============================================================================
If/Lsp Name/Svc-Id/RSVP-sess  State                Tx Intvl  Rx Intvl  Multipl
  Rem Addr/Info/SdpId:VcId    Protocols            Tx Pkts   Rx Pkts   Type
   LAG port                      LAG ID
```

```
-------------------------------------------------------------------------------
int-gre-tunnel               Up                   100       100       3
   10.0.0.2                  ospf2                1170      1169      central
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
*A:PE-1#
```

When BGP is configured instead of OSPF, the configuration of VPRN 1 on PE-1 is as follows:

```
*A:PE-1# configure
    service
        vprn 1 customer 1 create
            autonomous-system 64496
            route-distinguisher 64496:1
            vrf-target target:64496:1
            interface "int-gre-tunnel" tunnel create
                address 10.0.0.1/30
                bfd 100 receive 100 multiplier 3
                sap tunnel-1.private:1 create
                    ip-tunnel "gre-tunnel-1" create
                        dest-ip 10.0.0.2
                        gre-header
                        source 192.168.1.1
                        remote-ip 192.168.3.1
                        delivery-service 2
                        no shutdown
                    exit
                exit
            exit
            interface "loopback1" create
                address 172.31.1.1/32
                bfd 100 receive 100 multiplier 3
                loopback
            exit
            static-route-entry 172.16.3.1/32
                next-hop 10.0.0.2
                    no shutdown
                exit
            exit
            bgp
                group "group-1"
                    type internal
                    local-address 172.31.1.1
                    neighbor 172.16.3.1
                        bfd-enable
                    exit
                exit
                no shutdown
            exit
            no shutdown
        exit
```

The following command shows that the BFD session is up for protocol BGP on interface “int-gre-tunnel”:

```
*A:PE-1# show router 1 bfd session
===============================================================================
Legend:  wp = Working path   pp = Protecting path
===============================================================================
BFD Session
===============================================================================
If/Lsp Name/Svc-Id/RSVP-sess  State                Tx Intvl  Rx Intvl  Multipl
  Rem Addr/Info/SdpId:VcId    Protocols            Tx Pkts   Rx Pkts   Type
   LAG port                      LAG ID
-------------------------------------------------------------------------------
loopback1                     Up                   100       100       3
   172.16.3.1                 bgp                  3903      3902      central
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
*A:PE-1#
```

IP/GRE Termination – Advanced Topics

DSCP Value of Outer Delivery Header

- Default behavior — The DSCP value from the payload header is copied into the outer GRE header. This is a one to one copy and no QoS classifications are required. It is performed when no DSCP value is configured under the private gre-tunnel.
- Non default behavior — DSCP is configured under the private SAP (following example using af41).

```
*A:PE-1# configure
    service
        vprn 1 customer 1 create
            autonomous-system 64496
            route-distinguisher 64496:1
            vrf-target target:64496:1
            interface "int-gre-tunnel" tunnel create
                address 10.0.0.1/30
                bfd 100 receive 100 multiplier 3
                sap tunnel-1.private:1 create
                    ip-tunnel "gre-tunnel-1" create
                        dest-ip 10.0.0.2
                        gre-header
                        source 192.168.1.1
                        remote-ip 192.168.3.1
                        delivery-service 2
                        dscp af41
                        no shutdown
                    exit
---snip---
```

The log filter output: TOS=88 (DSCP=af41) in the public network.

```
*A:PE-1# show filter log 102

===============================================================================
Filter Log
===============================================================================
Admin state : Enabled
Description : (Not Specified)
Destination : Memory
Wrap        : Enabled
-------------------------------------------------------------------------------
Maximum entries configured : 1000
Number of entries logged   : 10
-------------------------------------------------------------------------------
2016/10/26 09:45:21  Ip Filter: 2:10  Desc:
SAP: tunnel-1.private:1 Direction: Egress  Action: Forward
Src MAC: 4a-c4-ff-00-02-c9  Dst MAC: 00-00-00-03-ff-9c  EtherType: 0800
Src IP: 10.0.0.1  Dst IP: 10.0.0.2  Flags: 0  TOS: 88  TTL: 64 Len: 84
Protocol: ICMP  Type: Echo Request  Code: 0
---snip---
```

IP-MTU

It is possible to configure the IP MTU of a private tunnel SAP interface. This sets the maximum IP packet size payload (including IP header) that can be sent into the tunnel (it applies to the packet size before the tunnel encapsulation is added).

```
*A:PE-1# configure
    service
        vprn 1 customer 1 create
            autonomous-system 64496
            route-distinguisher 64496:1
            vrf-target target:64496:1
            interface "int-gre-tunnel" tunnel create
                address 10.0.0.1/30
                ip-mtu 1476
            ---snip---
```

When an IPv4 packet needs to be forwarded to the tunnel and is larger than IP MT U bytes:

- If the DF bit is clear, the payload packet is IP fragmented to the MTU size prior to tunnel encapsulation.
- If the DF bit is set, the payload packet is discarded.

The IP-MTU range supported is from 512 to 9000 bytes.

The following command shows the configured IP MTU and the operational IP MTU for the GRE tunnel:

```
*A:PE-1# show router 1 interface "int-gre-tunnel" detail | match MTU
IP MTU           : 1476
IP Oper MTU      : 1476                 ICMP Mask Reply   : True
*A:PE-1#
```

Statistics and Accounting

Collect-stats can be configured under public and private SAPs.

For Public SAPs:

```
*A:PE-1# configure service ies 2 interface "int-tunnel-public" sap tunnel-
1.public:1 collect-stats
```

For Private SAPs:

```
*A:PE-1# configure service vprn 1 interface "int-gre-tunnel" sap tunnel-
1.private:1 collect-stats
```

Filtering, Policing, and QoS

An ip-filter and QoS policy can be applied to the ingress and egress traffic of the private and public SAPs.

Public SAPs:

```
*A:PE-1# configure
    service
        ies 2 customer 1 create
            interface "int-tunnel-public" create
                address 192.168.1.2/30
                tos-marking-state untrusted
                sap tunnel-1.public:1 create
                    ingress
                        qos 10
                        filter ip 1
                    exit
                    egress
                        qos 20
                        filter ip 1
                    exit
```

Private SAPs:

```
*A:PE-1# configure
    service
        vprn 1 customer 1 create
```

```
        route-distinguisher 64496:1
        vrf-target target:64496:1
        interface "int-gre-tunnel" tunnel create
            address 10.0.0.1/30
            sap tunnel-1.private:1 create
                ingress
                    qos 10
                    filter ip 1
                exit
                egress
                    qos 20
                    filter ip 1
                exit
            ---snip---
```

Mirroring

The public and private SAPs can be mirrored.

```
*A:PE-1# show debug
debug
    mirror-source 99
        sap tunnel-1.private:3 egress ingress
        sap tunnel-1.public:1 egress ingress
        no shutdown
    exit
exit
```

Conclusion

This chapter provides configuration and show commands for IP/GRE termination.

L2TP Network Server

This chapter provides information about L2TP network servers (LNS).

Topics in this chapter include:

Applicability

This feature is applicable to 7750 SR with at least one MS-ISA installed in an IOM3-XP (or later).

Initially, this chapter was written for SR OS release 11.0.R7, but the CLI in this edition is based on release 14.0.R2.

Overview

The Layer 2 Tunneling Protocol (L2TP) is a tunneling protocol commonly used to transport PPP sessions from an initiator known as an L2TP Access Concentrator (LAC) to an L2TP Network Server (LNS). L2TP is typically used for wholesaling residential broadband services. In this scenario, the LAC resides in the wholesaler's network and has a Layer-2 connection to an access concentrator such as a DSLAM. The LAC acts as the responder during the discovery phase (if PPPoE is used) and during PPP Link Control Protocol (LCP) negotiation. The LAC also performs an initial authentication of the subscriber. A successful authentication, typically from RADIUS, indicates to the LAC that PPP frames from this subscriber should be tunneled to an LNS at the indicated IP address. The LAC then tunnels the PPP frames from this subscriber over an L2TP tunnel to the LNS, where the PPP session is actually terminated. The terminology should be clear: PPP **sessions** are carried inside L2TP **tunnels**.

L2TP uses two types of messages; control messages and data messages. Control messages are used in the establishment, maintenance, and tearing down of tunnels and sessions. In order to provide extensibility and maximize interoperability, the payloads of control messages are encoded using Attribute Value Pairs (AVPs), some of which are applicable to all control messages, and some of which are specific to particular control messages. The L2TP header contains sequence number fields that must be present in control messages to allow for a reliable L2TP control channel that guarantees delivery. Data messages are used to encapsulate PPP frames being carried over the tunnel. Data messages are not retransmitted if packet loss occurs.

L2TP has a common fixed header format for both control and data messages, and a Type (T) bit in the header is used to indicate whether the packet is a control (1) or data (0) message. The L2TP packet is then carried in a transport protocol, and although the specifications allow L2TP to be directly encapsulated over Frame Relay, ATM, and UDP/IP, the latter is used almost exclusively.

The objective of this chapter is to provide a generic overview of how to configure the 7750 SR to support the LNS and LTS (L2TP Tunnel Switching) functions.

Example Topology

The simple example topology used through this chapter is shown in Figure 61. Both the LAC and the LNS participate in IS-IS and LDP, together with PE-1. All three devices form part of AS 64496 and peer using iBGP for the VPN-IPv4 address family. None of these protocols is mandatory for supporting LNS functionality; L2TP packets can ingress the system over any network interface as native IP or encapsulated as IP in MPLS, or through an IES/VPRN IP interface (SAP) as native IP. The MPLS data-plane within the example topology is chosen purely because of its simplicity and flexibility. Tester T1 simulates a DSLAM and one or more PPP clients, and is connected directly to the LAC. Although the LAC in this topology is a 7750 SR router, the configuration requirements of that device are beyond the scope of this example. Tester T2 provides a traffic source/sink capability and is connected directly to PE-1.

Figure 61 **Example Topology**

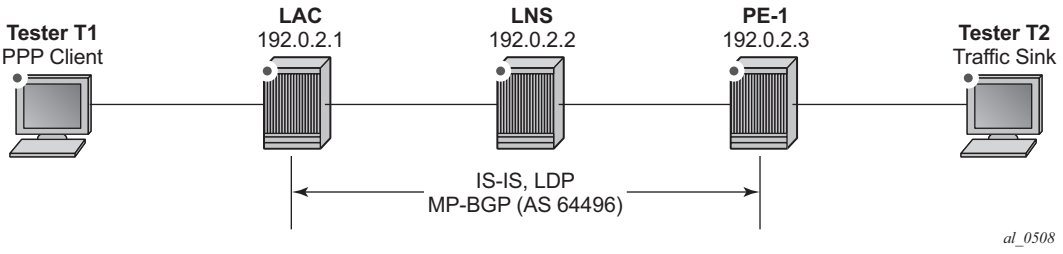

Hardware Configuration

To support LNS (and LTS) functionality at least one MS-ISA card is required, which must be configured as MDA type isa-bb, and must be housed in the carrier IOM, an IOM3-XP (or IOM3-XP-B/C). The MS-ISA performs L2TP data-plane encapsulation and de-capsulation, whereas the subscriber processing (Enhanced Subscriber Management or ESM) for PPP sessions is implemented within the carrier IOM.

```
configure
    card 1
        card-type iom3-xp
        mda 1
            mda-type m4-10gb-xp-xfp
            no shutdown
        exit
        mda 2
            mda-type isa-bb
            no shutdown
        exit
        no shutdown
    exit
exit
```

The MS-ISA is then configured to become a member of an **lns-group**. Up to six MS-ISAs can be configured to belong to one or more lns-groups. When two or more MS-ISAs belong to the same lns-group, by default PPP sessions are load-balanced over those MS-ISAs on a per-session basis.

```
configure
    isa
        lns-group 1 create
            mda 1/2
            no shutdown
        exit
    exit
exit
```

Configuration

ESM Base Configuration

For completeness, the following outputs contain the base ESM configuration that is applied to subscribers instantiated at the LNS throughout this chapter. Deviations from these base parameters are explicitly called out.

The SLA-Profile and Sub-Profile configurations have minimal parameters. The SLA-Profile uses the default ingress/egress QoS policy of 1, while the **no qos-marking-from-sap** command ensures that any subsequent marking is inherited from the egress QoS policy referenced in the SLA-profile, and not taken from the egress SAP. In order to do on-line accounting through RADIUS, the Sub-Profile also calls the relevant RADIUS accounting policy. Finally, the **sub-ident-policy** is configured with **use-direct-map-as-default** for the **sub-profile-map** and **sla-profile-map**, which means that the strings passed from RADIUS in the Vendor Specific Attributes (VSAs) **Alc-Subs-Prof-Str** and **Alc-SLA-Prof-Str** are interpreted verbatim so they are not used as string input to a mapping function.

```
configure
    subscriber-mgmt
        sla-profile "ESM-SLA-PROF" create
            egress
                no qos-marking-from-sap
            exit
        exit
        sub-profile "ESM-SUB-PROF" create
            radius-accounting
                policy "AAA-ACCT-POLICY"
            exit
        exit
        sub-ident-policy "all-subscribers" create
            sub-profile-map
                use-direct-map-as-default
            exit
            sla-profile-map
                use-direct-map-as-default
            exit
        exit
    exit
exit
```

Whilst it is entirely possible to authenticate subscribers locally using a local user database (LUDB), the more widely deployed approach is to use RADIUS, and this approach is therefore used implicitly throughout this chapter. The next output shows the **authentication-policy** AAA-AUTH-POLICY and **radius-accounting-policy** AAA-ACCT-POLICY. Both policies reference the **radius-server-policy** AAA-SUB-MGMT, which provides the context to configure the source-address to use for RADIUS messages and an associated routing context. The **radius-server-policy** then references a RADIUS server AAA-RADIUS, which in turn allows for configuration of the server IP address, the secret key to be used for message exchanges, and any other optional port configuration.

The intention is not to provide a complete description of all of the RADIUS parameters as this would distract from the objective of this chapter.

```
configure
    router
        radius-server
            server "AAA-RADIUS" address 172.16.1.1 secret vsecret1 create
```

```
                accept-coa
                pending-requests-limit 1024
            exit
        exit
    exit
    aaa
        radius-server-policy "AAA-SUB-MGMT" create
            servers
                router "Base"
                source-address 192.0.2.2
                server 1 name "AAA-RADIUS"
            exit
        exit
    exit
    subscriber-mgmt
        authentication-policy "AAA-AUTH-POLICY" create
            accept-authorization-change
            pppoe-access-method pap-chap
            include-radius-attribute
                nas-port-id
                nas-identifier
                access-loop-options
                calling-station-id remote-id
            exit
            radius-server-policy "AAA-SUB-MGMT"
        exit
        radius-accounting-policy "AAA-ACCT-POLICY" create
            no queue-instance-accounting
            session-accounting interim-update host-update
            update-interval 120
            include-radius-attribute
                circuit-id
                framed-ip-addr
                nas-identifier
                nas-port-id
                nas-port-type
                sla-profile
                sub-profile
                subscriber-id
                std-acct-attributes
            exit
            session-id-format number
            radius-server-policy "AAA-SUB-MGMT"
        exit
    exit
exit
```

Basic LNS Configuration

To illustrate the building blocks that are required to implement LNS functionality, a VPRN is used between the LAC and the LNS supporting an L2TP tunnel and terminating PPP sessions at the LNS. The required configuration for this VPRN at the LNS is shown in the following output. The unicast VPRN parameters such as **route-distinguisher** and **vrf-import**/**vrf-export** are not discussed here, only the parameters that are relevant to subscriber termination, which are equally applicable to VPRN and/or IES services.

The interface **loopback** represents a logical loopback interface which is used as the LNS endpoint address in L2TP signaling. The LAC has a corresponding interface with IP address 192.168.0.1. It also represents the unnumbered interface address referenced in the **subscriber-interface** context, meaning this IP address is used for the purpose of IPCP negotiation with incoming PPP sessions. Within the **subscriber-interface** context, the **group-interface** has a different definition than a conventional ESM **group-interface**. A conventional **group-interface** has one or more SAPs belonging to the same port or LAG. However, in the context of LNS, there are no SAPs. The group-interface also might terminate sessions within the same L2TP tunnel which are anchored on different MS-ISAs in a common lns-group. To accommodate this, the **group-interface** has the **creation-time** attribute **lns**. This attribute essentially means that the group-interface can terminate subscribers from more than one port/LAG; where port/LAG is interpreted as different MS-ISAs.

The **group-interface** then provides a **sap-parameters** context that allows for configuration of **sub-sla-mgmt** parameters that would typically be found under a SAP. These parameters are applicable to all subscribers terminated on this group-interface. In the example shown, only the **sub-ident-policy** is configured; therefore there is an expectation that other ESM parameters such as **sla-profile**, **sub-profile**, and **subscriber-id**, will be learned from a different source (in this example, they will be learned from RADIUS).

The static route blackholes prefix 10.48.127.0/24, ensuring this prefix is added to the route-table. Subscribers are allocated /32 addresses from this range, which must be advertised upstream to PE-1 to ensure end-to-end IP connectivity. This is facilitated through the **vrf-export policy** (not shown for conciseness).

Within the **l2tp** context, there exists a hierarchy of groups and tunnels. Groups reside directly under the **l2tp** context, and tunnels reside within the group context. Groups are intended to administratively organize tunnels that may share a common use or contain common parameters. The L2TP tunnel characteristics can be inherited from the group context, or overridden within the **tunnel** context. In the group context shown in the output below, the **lns-group 1** command refers to the **lns-group** previously configured at the ISA level. This is followed by the **local-address** command that indicates an IP address to be used as a source address for L2TP

signaling. The ppp context then defines the characteristics of how the PPP session will be established. In this case, the authentication mechanism is CHAP, and the previously configured RADIUS **authentication-policy** is used to authenticate the user. During the PPP session set-up, the LAC negotiates LCP and authentication parameters with the subscriber. Two AVPs, the **Proxy LCP** AVP and the **Proxy Authentication** AVP allow this information to be forwarded by the LAC to the LNS. This information can be accepted by the LNS, allowing PPP to continue with negotiation of IPCP, or it can be rejected, in which case the LNS initiates a new round of NCP and PPP authentication.The **proxy-authentication** and **proxy-lcp** commands allow the information contained in these AVPs to be accepted.

Finally, the **tunnel** context provides the context for explicit configuration related to this L2TP tunnel. The **peer** command indicates the far-end (LAC) IP address to which L2TP messages are addressed. The **password** is used to authenticate the far-end tunnel initiator, and is used in conjunction with the **challenge** parameter to implement a CHAP-like authentication mechanism. The default behavior is to never challenge the initiator (LAC); hence the **challenge always** setting is the inverse of this behavior. The **remote-name** is used to provide an additional level of security. When the Start Control Connection Request (SCCRQ) is received from the LAC to initiate the tunnel set-up it carries a mandatory **Host Name** AVP. The value of this AVP is compared to the name configured in the **remote-name**, and only tunnels with matching names are accepted. In a similar manner, the **local-name** parameter is used to populate the Host Name AVP sent by the LNS in the SCCRP, and can be used as a similar security feature at the LAC.

When two or more MS-ISAs belong to the same lns-group, by default PPP sessions are load-balanced over those MS-ISAs on a per-session basis. Although it is not shown in the configuration example below, it is worth highlighting that within each L2TP group context, an option exists to load-balance on a per-L2TP tunnel basis using the **load-balance-method per-tunnel** command. This can be useful, for example, when multiple sessions are received from a single subscriber (for example MLPPP member links) which must be handled within the same MS-ISA.

```
configure
    service
        vprn 1 customer 1 create
            vrf-import "vrf1-import"
            vrf-export "vrf1-export"
            route-distinguisher 64496:1
            auto-bind-tunnel
                resolution-filter
                    ldp
                exit
                resolution filter
            exit
            interface "system" create
                address 192.168.0.2/32
                loopback
            exit
            subscriber-interface "LNS-SUB-INT" create
```

```
                    unnumbered 192.168.0.2
                    group-interface "LNS-GROUP-INT" lns create
                        sap-parameters
                            sub-sla-mgmt
                                sub-ident-policy "all-subscribers"
                            exit
                        exit
                    exit
                    static-route-entry 10.48.127.0/24
                        black-hole
                            no shutdown
                        exit
                    exit
                    l2tp
                        group "L2TP-GROUP-1" create
                            hello-interval 60
                            idle-timeout 600
                            lns-group 1
                            local-address 192.168.0.2
                            ppp
                                authentication chap
                                authentication-policy "AAA-AUTH-POLICY"
                                default-group-interface "LNS-GROUP-INT" service-id 1
                                keepalive 10 hold-up-multiplier 3
                                proxy-authentication
                                proxy-lcp
                            exit
                            tunnel "L2TP-TUNNEL-1" create
                                challenge always
                                local-name "LNS"
                                peer 192.168.0.1
                                remote-name "LAC"
                                password tunnelpwd
                                no shutdown
                            exit
                            no shutdown
                        exit
                        no shutdown
                    exit
                    no shutdown
                exit
            exit
        exit
```

As previously described, RADIUS is used to authenticate the subscriber, which upon successful authentication returns the ESM parameters, Subscriber-ID (**Alc-Subsc-ID-Str**), SLA-Profile (**Alc-SLA-Prof-Str**), and Sub-Profile (**Alc-Subsc-Prof-Str**) as needed for instantiating the subscriber in SR OS. These parameters could be obtained locally on the LNS using the **def-sub-id**, **def-sla-profile** and **def-sub-profile** commands under the **group-interface sap-parameters**. This enables a mechanism to provide default parameters in the absence of obtaining them from another source. However, passing them from RADIUS has some benefits, such as:

- It is comparatively easy to provide different SLA- and Sub-Profiles to different users, which can be used to differentiate service levels.

- If a RADIUS infrastructure is in place and used to provide ESM parameters, it is comparatively easy to extend that infrastructure to provide for mid-session changes of those parameters (such as **sla-profile** and **sub-profile**) using a Change of Authorization (CoA).

The next output provides an example of a RADIUS users file entry for the test subscriber. In addition to the afore-mentioned ESM parameters, the Alc-Serv-ID VSA is used to indicate the service number in which this subscriber must be terminated (in this case VPRN 1 as previously configured), while the Alc-Interface VSA is used to indicate the relevant group-interface within that service. If it is intended that all PPP sessions ingressing on a particular L2TP group are all to be terminated within a common service and group-interface, it is not necessary to send the Alc-Serv-ID and Alc-Interface VSAs from RADIUS to indicate the service and group-interface, but rather a default service and group-interface can be specified within the **ppp** context of the l2tp group using the parameter **default-group-interface** <*name*> **service-id** <*number*>. The remainder of the attributes in the output are well-known standard attributes.

```
subscriber1@isp.net     Cleartext-Password := "letmein"
                        Alc-Subsc-ID-Str = "subscriber1@isp.net",
                        Alc-Subsc-Prof-Str = "ESM-SUB-PROF",
                        Alc-SLA-Prof-Str = "ESM-SLA-PROF",
                        Alc-Serv-Id = "1",
                        Alc-Interface = "LNS-GROUP-INT",
                        Service-Type = Framed-User,
                        Framed-Protocol = PPP,
                        Framed-IP-Address = 10.48.127.27,
```

L2TP Tunnel Set-up

Before the PPP session can be terminated at the LNS, an L2TP tunnel must be established between the LAC and LNS. This is achieved using a three-way control message exchange of Start-Control-Connection-Request (SCCRQ), Start-Control-Connection-Reply (SCCRP), and Start-Control-Connection-Connected (SCCN). All of these messages are explicitly acknowledged by the peer using the sequence numbers (number sent, number received) in the L2TP header, thereby creating a reliable control channel. The acknowledgment can be piggy-backed in a corresponding control message, or it can be an explicit acknowledgment using a control packet with only an L2TP header, known as a Zero-Length Body (ZLB) message.

The SCCRQ is used to initialize the tunnel between LAC and LNS, and although it can be sent by either the LAC or LNS, it is typically sent by the LAC toward the LNS (as in this example). The SCCRQ contains a number of mandatory AVPs, denoted by the M-bit in the AVP header (set to 1), including Message Type, Protocol Version, Host Name, Framing Capabilities, and Assigned Tunnel ID. It can also contain a number of optional AVPs, such as Vendor Name, and Firmware Revision, which can be ignored by the recipient if they are unrecognized.

```
1 2016/05/27 07:57:17.31 UTC MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.0.1:1701 -> 192.168.0.2:1701
tunnel 0 session 0, ns 0 nr 0, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        StartControlConnectionRequest(1)
    AVP ProtocolVersion(0,2), flags: mandatory, reserved=0
        version=1, revision=0
    AVP HostName(0,7), flags: mandatory, reserved=0
        "LAC"
    AVP WindowSize(0,10), flags: mandatory, reserved=0
        64
    AVP FramingCapabilities(0,3), flags: mandatory, reserved=0
        sync=no, async=no
    AVP BearerCapabilities(0,4), flags: mandatory, reserved=0
        digital=yes, analogue=no
    AVP FirmwareRevision(0,6), flags:, reserved=0
        3584
    AVP VendorName(0,8), flags:, reserved=0
        "Nokia"
    AVP AssignedTunnelId(0,9), flags: mandatory, reserved=0
        1796"
```

The SCCRP is sent in response to the SCCRQ and is used to indicate that the parameters in the SCCRQ were acceptable and that the establishment of the L2TP tunnel can continue. The SCCRP contains the same mandatory AVPs and can contain the same optional AVPs as the SCCRQ, but an additional optional AVP is the Challenge AVP which is included as a result of the **challenge always** and **password** parameters configured within the **tunnel** context.

```
2 2016/05/27 07:57:17.31 UTC MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.168.0.2:1701 -> 192.168.0.1:1701
tunnel 1796 session 0, ns 0 nr 1, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        StartControlConnectionReply(2)
    AVP ProtocolVersion(0,2), flags: mandatory, reserved=0
        version=1, revision=0
    AVP HostName(0,7), flags: mandatory, reserved=0
        "LNS"
    AVP WindowSize(0,10), flags: mandatory, reserved=0
        64
    AVP FramingCapabilities(0,3), flags: mandatory, reserved=0
        sync=no, async=no
    AVP BearerCapabilities(0,4), flags: mandatory, reserved=0
        digital=yes, analogue=no
    AVP FirmwareRevision(0,6), flags:, reserved=0
        3584
    AVP VendorName(0,8), flags:, reserved=0
```

```
        "Nokia"
    AVP AssignedTunnelId(0,9), flags: mandatory, reserved=0
        1146
    AVP Challenge(0,11), flags: mandatory, reserved=0
         45 03 b9 a0 c6 d9 eb 83 bc 8e e3 a1 9d 91 35 6b
         99 8f 29 5e 18 db c3 8e aa 50 90 08 d9 c7 68 e0
         df b2 d2 41 b7 ff 54 63 70 "
```

The response to the SCCRP, and the completion of the three-way message exchange is the SCCN. The only mandatory AVP for the SCCN is the Message Type, and since the SCCRP contained a Challenge AVP, the SCCN also contains an AVP Challenge Response. If this response is not satisfactory to the LNS, it generates a Stop-Control-Connection-Notification (StopCCN) with a result code indicating that the requester is not authorized, and subsequently removes any associated tunnel state.

```
3 2016/05/27 07:57:17.30 UTC MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.0.1:1701 -> 192.168.0.2:1701
tunnel 1146 session 0, ns 1 nr 1, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        StartControlConnectionConnected(3)
    AVP ChallengeResponse(0,13), flags: mandatory, reserved=0
         ef b8 1e 7f 84 e6 64 02 10 54 ad eb 00 57 dd 82 "
```

With a successful three-way exchange completed, the L2TP tunnel is established. A snapshot view of all L2TP tunnels within the relevant routing context can be displayed using the command in the following output. The Loc-Tu-ID and Rem-Tu-ID are the local and remote tunnel IDs passed in the Assigned Tunnel Id AVP in the SCCRP and SCCRQ respectively. The Conn ID, or connection Id, is a locally significant parameter used for the purpose of identifying a particular tunnel, and is a 32-bit representation of the local tunnel Id (1146 * 65536 = 75104256). It is the connection ID that is used, for example, in event log entries for this tunnel. The state is shown as **established**, because one or more PPP sessions are running over the tunnel. The state can also be **establishedIdle** meaning that although the tunnel is up and established, there are no PPP sessions active within the tunnel.

```
*A:LNS# show router 1 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State              Blacklist-state   Ses Active
  Group                                                             Ses Total
    Assignment
-------------------------------------------------------------------------------
75104256   1146      1796      established        not-blacklisted   1
  L2TP-GROUP-
1                                                      1
    L2TP-TUNNEL-1
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
*A:LNS#
```

Once a tunnel is established, maintenance and health-checking of that tunnel is achieved using a keepalive mechanism that employs Hello control messages. The Hello message contains only one AVP, the **Message Type** AVP, which indicates it is a Hello message. The Hello messages operate asynchronously between the peers. There is no echo request and echo response function, but simply a Hello followed by an acknowledgment. The Hello is acknowledged in the same way as other control messages, using either piggy-backing or ZLB acknowledgments. This asynchronous behavior allows for either end of the tunnel to be configured for different Hello intervals (they are not negotiated), or even for one end not send Hellos at all.

```
20 2016/05/27 07:58:23.12 UTC MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.168.0.2:1701 -> 192.168.0.1:1701
tunnel 1796 session 0, ns 2 nr 4, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        Hello(6)"
21 2016/05/27 07:58:23.12 UTC MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.0.1:1701 -> 192.168.0.2:1701
tunnel 1146 session 0, ns 4 nr 3, flags:, reserved=0"
```

The Hello interval at the LNS is configurable under the l2tp, group, or tunnel contexts using the **hello-interval** parameter. The range is 60 to 3600 seconds, with the default being 300 seconds. The **hello-interval infinite** option suppresses sending of Hellos. If the system sends a Hello message and does not get an acknowledgment, it will re-transmit the Hello message as many times as the value of the **max-retries-estab** parameter, each time with an exponential(Note:) back-off. The **max-retries-estab** parameter can be configured in the l2tp, group, or tunnel contexts. The default value is 5, and if no acknowledgment is received before this value is exceeded, the tunnel is declared down and a StopCCN is sent toward the peer.

Note: (1) The retry interval starts with 1 second and doubles on each retry with a maximum-interval of 8 seconds. For example, using a max-retries-estab value of 7 results in a retry of [1, 2, 4, 8, 8, 8 seconds]

The StopCCN is a message that can be generated by either LAC or LNS and is used to inform its peer that the tunnel is being closed. This implicitly means that all PPP sessions carried within that tunnel are also being closed without any associated control messages for those sessions. The StopCCN must contain the **Message Type** and **Tunnel ID** AVPs, and additionally carries a **Result Code** AVP with result code and error code fields to indicate to the peer the reason for the tunnel closure.

```
36 2016/05/27 08:03:33.40 UTC MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.0.1:1701 -> 192.168.0.2:1701
tunnel 1146 session 0, ns 6 nr 8, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        StopControlConnectionNotification(4)
    AVP ResultCode(0,1), flags: mandatory, reserved=0
        result-code: "generalRequestToClearControlConnection"(1), error-code: "n
```

```
oGeneralError"(0)
        error-msg: "operator request"
    AVP AssignedTunnelId(0,9), flags: mandatory, reserved=0
        1796"
```

The tunnel **Connection Id** can be used as an additional argument to display the details of a particular tunnel when multiple tunnels are present. The following output is an example of this taken just after the L2TP tunnel has been closed by the LAC peer, and is intentionally taken at this time to illustrate the purpose of some of the fields shown in the output. The State is moved to **closedByPeer**, and the Stop CCN Result field and Error Message field respectively contain the result code and error code of the Result Code AVP received from the LAC in the StopCCN. Because the tunnel is now in a closedByPeer state, all state and information related to this tunnel is removed from the system after a period defined by the Destruct Timeout (shown in the output as Destruct TO). The intention of the Destruct Timeout is to retain information about the tunnel closure which might aid operational communities. The default value as shown is 60 seconds, but it can be configured using the **destruct-timeout** parameter in the l2tp, group, or tunnel contexts. The remainder of the fields in the output are the operational parameters of the tunnel and are self-explanatory.

```
*A:LNS# show router 1 l2tp tunnel connection-id 75104256 detail

===============================================================================
L2TP Tunnel Status
===============================================================================

Connection ID: 75104256
State        : closedByPeer
IP           : 192.168.0.2
UDP          : 1701
Peer IP      : 192.168.0.1
Peer UDP     : 1701
Tx dst-IP    : 192.168.0.1
Tx dst-UDP   : 1701
Rx src-IP    : 192.168.0.1
Rx src-UDP   : 1701
Name         : LNS
Remote Name  : LAC
Assignment ID: L2TP-TUNNEL-1
Group Name   : L2TP-GROUP-1
Acct. Policy : N/A
Error Message: operator request

                                       Remote Conn ID    : 117702656
Tunnel ID         : 1146               Remote Tunnel ID  : 1796
Preference        : 50                 Receive Window    : 64
Hello Interval (s): 60
Idle TO (s)       : 600                Destruct TO (s)   : 60
Max Retr Estab    : 5                  Max Retr Not Estab: 5
Session Limit     : 32767              AVP Hiding        : never
Transport Type    : udpIp              Challenge         : always
Time Started      : 05/27/2016 07:57:17 Time Idle         : 05/27/2016 08:02:23
Time Established  : 05/27/2016 07:57:17 Time Closed       : 05/27/2016 08:03:33
Stop CCN Result   : generalReq         General Error     : noError
```

```
Blacklist-state   : not-blacklisted
Set Dont Fragment : true

Failover
State             : not-recoverable
Recovery Conn ID  : N/A
Recovery state    : not-applicable
Recovered Conn ID : N/A
Recovery method   : mcs
Track SRRP        : (Not specified)
Ctrl msg behavior : handle
Recovery time (ms)
Requested         : N/A
Peer              : N/A
-------------------------------------------------------------------------------
===============================================================================
*A:LNS#
```

PPP Session Set-up

Once the L2TP tunnel is in place, the process of establishing a PPP session can commence. Once again there is a three-way control message exchange used for establishing a session within an L2TP tunnel, consisting of the Incoming-Call-Request (ICRQ), Incoming-Call-Reply (ICRP), and Incoming-Call-Connected (ICCN). Given that they are control messages, they are all explicitly acknowledged using piggybacking or ZLB acknowledgments.

The ICRQ is sent from the LAC to the LNS to indicate that it has received an incoming call (PPP session) and that a session needs to be established between the two peers for this call. The ICRQ provides enough information about the call for the LNS to make a decision about whether it should answer the call or not. The ICRQ contains the Message Type and Assigned Session ID AVPs as well as a Call Serial Number AVP, which can be used on both the LAC and LNS as an easy reference to the call for troubleshooting purposes. The ICRQ can also carry optional AVPs including Calling Number and Access Line Information AVPs (RFC 5515) such as Circuit ID, Remote ID, Actual Data Rate Upstream, and Actual Data Rate Downstream.

```
5 2016/05/27 07:57:17.31 UTC MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.0.1:1701 -> 192.168.0.2:1701
tunnel 1146 session 0, ns 2 nr 1, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallRequest(10)
    AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
        9646
    AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
        15672
    AVP CallingNumber(0,22), flags: mandatory, reserved=0
        "BBEU4966723450"
    AVP AgentCircuitId(3561,1), flags:, reserved=0
        "dslam142-atm4/2/7:0.101"
    AVP AgentRemoteId(3561,2), flags:, reserved=0
```

```
        "BBEU4966723450"
    AVP ActDataRateUp(3561,129), flags:, reserved=0
        2048000
    AVP ActDataRateDown(3561,130), flags:, reserved=0
        8192000"
```

The ICRP is sent by the LNS toward the LAC in response to the ICRQ to indicate that the parameters in the ICRQ were acceptable, and that the LAC should go ahead and proceed with the call. The ICRP contains only two AVPs; the Message Type and the Assigned Session ID. The Assigned Session ID values are local to each peer as opposed to a negotiated or agreed-upon value.

```
7 2016/05/27 07:57:17.31 UTC MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.168.0.2:1701 -> 192.168.0.1:1701
tunnel 1796 session 9646, ns 1 nr 3, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallReply(11)
    AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
        26192"
```

The final message in the three-way exchange used for establishing sessions within the tunnel is the ICCN. It is sent by the LAC to the LNS to indicate that the call has been answered, so the L2TP session is moved to the **established** state. It also provides additional information on parameters that were used to answer the call which may not have been available when the ICRQ was sent (although it is likely that in most cases they were available). At a minimum, the ICCN must contain the Message Type, Framing Type and TX Connect Speed AVPs. The TX Connect Speed defines the speed in bits-per-second from the perspective of traffic flowing from the LAC toward the subscriber (i.e. the LAC downstream rate) and, for best accuracy, can be derived by the LAC from the PPP Broadband Forum Access Line Characteristic tags inserted by the access node (Appendix C TR-101). The TX Connect Speed can be useful for indirect setting of a Hierarchical QoS (H-QoS) aggregate rate. It is indirect because the LNS cannot infer and set an aggregate rate based directly on the TX Connect Speed AVP, but rather the TX Connect Speed is passed to RADIUS (using the **include-radius-attribute access-loop-option** parameter in the authentication-policy), which in turn may pass the aggregate rate to the LNS in a QoS override VSA. This is discussed further in the QoS section.

A number of optional AVPs can also be present providing information from the LCP negotiation between the LAC and client. These include Initial Receive, Last Transmit and Last Receive LCP Config Requests, together with Proxy Authentication Type, Name, Challenge and Response. These parameters allow the LNS to either force a renegotiation of LCP, or to continue with the PPP session and move onto the IPCP phase. The final AVP present in the ICCN shown is the RX Connect Speed AVP, which is the opposite of the TX Connect Speed and defines the speed in bits-per-second from the perspective of traffic flowing from the subscriber toward the LAC.

```
9 2016/05/27 07:57:17.31 UTC MINOR: DEBUG #2001 vprn1 L2TP(v2, ctrl, ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.0.1:1701 -> 192.168.0.2:1701
```

```
tunnel 1146 session 26192, ns 3 nr 2, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallConnected(12)
    AVP FramingType(0,19), flags: mandatory, reserved=0
        sync=no, async=no
    AVP TxConnectSpeed(0,24), flags: mandatory, reserved=0
        4294967295
    AVP InitialRxLcpConfReq(0,26), flags:, reserved=0
         01 04 05 d4
         [1] MRU: 1492
    AVP LastTxLcpConfReq(0,27), flags:, reserved=0
         01 04 05 d4 03 05 c2 23 05 05 06 3c 2e bd 3b
         [1] MRU: 1492
         [3] Authentication-Protocol: 0xc223 (CHAP), Algorithm = 5 (MD5)
         [5] Magic-Number: 0x3c2ebd3b
    AVP LastRxLcpConfReq(0,28), flags:, reserved=0
         01 04 05 d4
         [1] MRU: 1492
    AVP ProxyAuthenType(0,29), flags:, reserved=0
        chap(2)
    AVP ProxyAuthenName(0,30), flags:, reserved=0
         "subscriber1@isp.net"
    AVP ProxyAuthenChallenge(0,31), flags:, reserved=0
         cb 33 7e 46 fd 62 87 5d f8 cb a7 e9 a3 94 74 21
         75 31 d4 81 e4 2c 1d 0f b2 1e 11 b6 32 36 ea 7d
         69 69 43 e6 4b 4b 44 c3 17 6c 2d 3b 01 a1 db 76
         53 31 77
    AVP ProxyAuthenId(0,32), flags:, reserved=0
        id=1, reserved=0
    AVP ProxyAuthenResponse(0,33), flags:, reserved=0
         da 56 92 c3 81 06 a1 e3 87 e5 19 f1 9c 30 3f 71
    AVP RxConnectSpeed(0,38), flags:, reserved=0
        4294967295"
```

On completion of the three-way control message exchange required for session set-up, the LNS authenticates the user in the incoming call. In this example RADIUS is used, which returns the standard and Vendor-Specific attributes previously defined in the users file. A successful authentication allows the LNS to move to the IPCP phase with the subscriber. In this example, RADIUS returns the IP address in the standard attribute Framed-IP-Address, but equally local pooling with a DHCP server could be used. For conciseness, the IPCP phase is not detailed within this example because the process is reasonably well-known and understood. However, on completion of IPCP the subscriber is instantiated and the L2TP session becomes active. The Tunnel-ID and Session-ID parameters are locally generated numbers that are passed in L2TP control messages. As previously described, the Connection Id is a locally significant parameter that is a 32-bit representation of the local tunnel Id (1146 * 65536 = 751304256). The ID field is again a locally significant parameter used to identify the L2TP session, and is again represented as a 32-bit number. It is derived from a sum of the Control Connection ID plus the Session ID (751304256 + 26192 = 75130448).

```
*A:LNS# show router 1 l2tp session

===============================================================================
```

```
L2TP Session Summary
===============================================================================
ID                  Control Conn ID     Tunnel-ID   Session-ID  State
-------------------------------------------------------------------------------
75130448            75104256            1146        26192       established
  subscriber1@isp.net
  interface: LNS-GROUP-INT
  service-id: 1
  10.48.127.27
-------------------------------------------------------------------------------
No. of sessions: 1
===============================================================================
*A:LNS#
```

The PPP session is also recorded in the subscriber-host table of VPRN 1 and a forwarding state of **Fwding** indicates that all attributes and resources associated with this subscriber are correctly installed and activated within the system. The subscriber username is shown, as is its MAC address and IP address. The IP address has an origin of IPCP. The fact that a MAC address is displayed here is somewhat misleading because this is a PPP over L2TP session, which does not have a MAC address present in any of its headers. When the MS-ISA removes the L2TP header it converts the PPP packet to PPPoE for ease of subsequent processing. As a result of this, the MS-ISA generates a fake MAC address, and this is the MAC address shown. The displayed SAP 1/2/lns-esm:1.259 is automatically generated by the system. Each operational MS-ISA that is part of the lns-group creates two internal objects, known as lns-net and lns-esm. These objects essentially represent a network-side (lns-net) and subscriber-side (lns-esm) of each MS-ISA.

When the first L2TP session within this service is established the system creates one lns-esm SAP where the first two digits indicate the MDA slot (1/2) where the MS-ISA is installed, and the last two numbers are the internal Q-in-Q tags used through the MS-ISA (1.259) (Note:). If there are more than one MS-ISA active in the lns-group, a second session would be load-balanced onto this MS-ISA, and a second lns-esm SAP would be created, until a maximum of six SAPs is reached, which represents the maximum number of supported active MS-ISA boards.

Note: (2) The internal Q-in-Q tag value is of little relevance, but for informational purposes is derived from the group-interface If index.

```
*A:LNS# show service id 1 subscriber-hosts

===============================================================
Subscriber Host table
===============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID Origin       Fwding State
---------------------------------------------------------------
[1/2/lns-esm:1.259]    subscriber1@isp.net
```

```
  10.48.127.27
    00:00:04:7a:66:50    1         IPCP         Fwding
------------------------------------------------------------
Number of subscriber hosts : 1
============================================================
*A:LNS#
```

It is also possible to view the internal lns-net object, shown in the next output as interface name **_tmnx_lns-in-1/2** with port number **1/2/lns-net:1***. To further clarify (and reiterate), the lns-esm and lns-net are simply internal objects used to route L2TP traffic through the MS-ISA board. Upstream traffic (subscriber to LNS) ingresses through lns-net into the MS-ISA where the L2TP header is decapsulated before PPP packets are presented to the service group-interface through lns-esm. Downstream traffic (LNS to subscriber) passes through lns-esm into the MS-ISA where the PPP packets are encapsulated in L2TP before egressing through lns-net and being routed toward the destination.

```
*A:LNS# show service id 1 interface

===============================================================================
Interface Table
===============================================================================
Interface-Name                   Adm          Opr(v4/v6)  Type    Port/SapId
   IP-Address                                              PfxState
-------------------------------------------------------------------------------
loopback                         Up           Up/Down     VPRN    loopback
   192.168.0.2/32                                                  n/a
LNS-SUB-INT                      Up           Up/Down     VPRN S* subscriber
   Unnumbered If[192.168.0.2]                                      n/a
LNS-GROUP-INT                    Up           Up/Down     VPRN G* bbg-5.lns-esm
_tmnx_lns-in-1/2                 Up           Up/Down     VPRN    1/2/lns-net:1*
   -                                                           -
-------------------------------------------------------------------------------
Interfaces : 4
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:LNS#
```

Wholesale/Retail

In the example configuration used so far the L2TP tunnel/session is terminated in VPRN 1, and the subscriber is also terminated in the same VPRN 1. However, a common requirement is to build per-customer VRFs (VPRNs), particularly for business users. To meet this requirement, the so-called ‘Wholesale/Retail’ model is used, which provides a mechanism to terminate the subscriber in a different service context from the service which actually terminated the L2TP tunnel/session.

To achieve this, a second service is created which becomes the ‘Retail VRF’, or customer-specific VRF, and the previously defined VPRN 1 becomes the Wholesale VRF (which actually requires no further configuration). The necessary configuration for the Retail VRF is as follows and its parameters have been previously explained. Although they may seem obvious, there are a couple of points that are worth revisiting. The **vrf-import** and **vrf-export** parameters are used to reference policies to import/export VPN-IPv4/v6 prefixes with the customer-specific Route-Target Extended Communities. Given that a different routing context and unique Route-Targets are used for this Retail VRF, it is perfectly feasible to re-use the same IP address in VPRN 2 as was used in VPRN 1 for the unnumbered subscriber-interface. The group-interface has a different name from the group-interface in VPRN 1, but this is simply for illustration purposes and both group-interfaces can have the same name if a standard naming convention is required. More importantly, the group-interface must have the creation-time attribute **lns** to allow subscriber termination without SAPs. The static route blackholes prefix 10.10.148.0/24, ensuring this prefix is added to the route-table. This IP address range is used to allocate addresses to subscribers, and is therefore advertised in VPN-IPv4.

```
configure
    service
        vprn 2 customer 1 create
            vrf-import "vrf2-import"
            vrf-export "vrf2-export"
            route-distinguisher 64496:2
            auto-bind-tunnel
                resolution-filter
                    ldp
                exit
                resolution filter
            exit
            interface "loopback" create
                address 192.168.0.2/32
                loopback
            exit
            subscriber-interface "VPRN2-SUB-INT" create
                unnumbered 192.168.0.2
                group-interface "VPRN2-GROUP-INT" lns create
                    sap-parameters
                        sub-sla-mgmt
                            sub-ident-policy "all-subscribers"
                        exit
                    exit
                exit
            exit
            static-route-entry 10.10.48.0/24
                black-hole
                    no shutdown
                exit
            exit
            no shutdown
        exit
    exit
exit
```

In addition to the Retail VRF configuration, the RADIUS entry for the subscriber returns **Alc-Serv-Id** VSA with a value of 2 to indicate the Retail VRF Service Id, while the **Alc-Interface** VSA refers to the group-interface name within that Retail VRF.

```
subscriber2@isp.net     Cleartext-Password := "letmein"
                        Alc-Subsc-ID-Str = "subscriber2@isp.net",
                        Alc-Subsc-Prof-Str = "ESM-SUB-PROF",
                        Alc-SLA-Prof-Str = "ESM-SLA-PROF",
                        Alc-Serv-Id = "2",
                        Alc-Interface = "VPRN2-GROUP-INT",
                        Service-Type = Framed-User,
                        Framed-Protocol = PPP,
                        Framed-IP-Address = 10.10.148.22
```

In this Wholesale/Retail scenario, the high-level functions are as follows:

- The L2TP tunnel and session are terminated in the Wholesale VRF (in this example, VPRN 1).
- When the LNS receives the ICCN for the session, it authenticates the user (in this example via RADIUS).

RADIUS returns the Retail VRF service Id and group-interface. If RADIUS returns IP address information this address is used for the purpose of IPCP negotiation with the subscriber within the Retail VRF (in this example, VPRN 2). If RADIUS does not return IP address information, it can be derived from either of the following:

- A DHCP client function within the group-interface, which is used to obtain an IP address from a local or remote DHCP server.
- The local-address-assignment feature, which directly accesses a local DHCP server through an internal procedure call (the server pool name must be learned through RADIUS, LUDB, or default-pool-name).

Once the subscriber is activated, the PPP session and subscriber-host can be seen in VPRN 2. The description field of the **show service id 2 ppp session** command is however somewhat misleading. It is automatically concatenated from the VPRN that terminated the L2TP tunnel, the tunnel Connection Id, the local tunnel Id, and the L2TP session Id. It should not be misinterpreted as meaning that the subscriber has been terminated in VPRN 1.

```
*A:LNS# show service id 2 ppp session

===============================================================================
PPP sessions for service 2
===============================================================================
User-Name
  Descr.
          Up Time       Type  Termination     IP/L2TP-Id/Interface-Id MC-Stdby
-------------------------------------------------------------------------------
subscriber2@isp.net
  vprn:1 connid:281635894 tid:4297 sid:27702
```

```
          0d 00:00:19   oL2tp local           10.10.148.22
-------------------------------------------------------------------------------
No. of PPP sessions: 1
===============================================================================
*A:LNS#
```

QoS

To this point the subscriber PPP sessions terminated by the LNS have been instantiated using the default SAP-ingress/egress QoS policies (policy 1), with a single queue and no use of H-QoS. This section demonstrates the use of slightly more complex QoS policies that employ H-QoS, with the intention of providing an overview of those capabilities.

For subscriber termination in broadband networks, it is fairly commonplace to use one or more policers on ingress, and not apply an aggregate rate-limit on ingress (upstream) traffic. Whilst this is possible in SR OS for general ESM subscriber termination, policers are not supported when the system is functioning as an LNS. It is therefore necessary to use one or more queues on ingress with the usual considerations with regard to the use of service-queuing or shared-queuing. Conversely, on egress (downstream) it is common to see more than one queue in use for different services, particularly for business services, with an aggregate rate applied to the subscriber through the use of H-QoS. For example, assume that there are three classes in use; Best-Effort (BE), Assured-Forwarding (AF), and Expedited Forwarding (EF). This section will look at two ways to achieve this. Firstly using a conventional H-QoS scheduler, and secondly using an egress Port-Scheduler.

The SAP-ingress QoS policy classifies traffic into three Forwarding Classes (FCs) and maps those FCs to a single queue. Ingress traffic is not rate-limited (default PIR in queue 1 is max), and queue 1 is not mapped to a parent H-QoS scheduler.

```
configure
    qos
        sap-ingress 10 create
            queue 1 create
            exit
            queue 11 multipoint create
            exit
            fc "af" create
                queue 1
            exit
            fc "be" create
                queue 1
            exit
            fc "ef" create
                queue 1
            exit
            dscp be fc "be"
            dscp ef fc "ef"
```

```
            dscp af31 fc "af"
        exit
    exit
exit
```

A scheduler policy is created having a single a tier 1 scheduler with a rate-limit of 8Mb/s.

```
configure
    qos
        scheduler-policy "Subscriber-Aggregate-Policy" create
            tier 1
                scheduler "Aggregrate-Rate" create
                    rate 8000
                exit
            exit
        exit
    exit
exit
```

The SAP-egress QoS policy performs egress classification and maps classified traffic to the relevant FC, which in turn is mapped to its own queue. All queues are mapped to the previously configured tier 1 scheduler **Aggregate-Rate** such that queue 3 (EF) is allocated bandwidth first, and queue 1 (BE) and 2 (AF) are allocated bandwidth next in a 1:4 ratio.

```
configure
    qos
        sap-egress 10 create
            queue 1 create
                parent "Aggregrate-Rate" level 2 weight 20
            exit
            queue 2 best-effort create
                parent "Aggregrate-Rate" level 2 weight 80
            exit
            queue 3 expedite create
                parent "Aggregrate-Rate" cir-level 3
                rate 1024 cir 1024
            exit
            fc af create
                queue 2
            exit
            fc be create
                queue 1
            exit
            fc ef create
                queue 3
            exit
            dscp be fc "be"
            dscp ef fc "ef"
            dscp af31 fc "af"
        exit
    exit
exit
```

To this point, the QoS configuration is no different from a typical SAP-level QoS application. To make it applicable to ESM, the previously configured SAP-ingress and SAP-egress QoS policies must be referenced in the ingress/egress contexts of the sla-profile, respectively. Equally, the H-QoS scheduler-policy must be referenced in the ingress/egress contexts of the sub-profile. In this example, H-QoS is only used on egress, and as a result the scheduler-policy is referenced only in the egress context.

```
configure
    subscriber-mgmt
        sla-profile "ESM-SLA-PROF" create
            ingress
                qos 10
                exit
            exit
            egress
                qos 10
                exit
                no qos-marking-from-sap
            exit
        exit
        sub-profile "ESM-SUB-PROF" create
            collect-stats
            radius-accounting
                policy "AAA-ACCT-POLICY"
            exit
            egress
                scheduler-policy "Subscriber-Aggregate-Policy"
                exit
            exit
        exit
    exit
exit
```

The queues assigned to the subscriber through the above SAP-ingress/egress QoS policies, together with accumulative statistics can be viewed using the **show service active-subscribers subscriber** <*name*> **detail** command (real time rates can be seen using the **monitor** command). The H-QoS scheduler hierarchy, with the SAP-egress queues mapped as child queues to a parent scheduler can be validated using the command **show qos scheduler-hierarchy subscriber** <*name*> **egress**. The **detail** argument as an extension of this command provides a significant amount of detail on real-time bandwidth allocated to each queue by the scheduler in the within-CIR and above-CIR passes. It also provides a useful snapshot on offered traffic loads in Kb/s on a per-queue basis.

```
*A:LNS# show qos scheduler-hierarchy subscriber "subscriber2@isp.net" egress

===============================================================================
Scheduler Hierarchy - Subscriber subscriber2@isp.net
===============================================================================
Egress Scheduler Policy : Subscriber-Aggregate-Policy
-------------------------------------------------------------------------------
Root (Egr)
```

```
| slot(1)
|--(S) : Aggregrate-Rate
|   |
|   |--(Q) : Sub=subscriber2@isp.net:ESM-SLA-PROF 2->1/2/lns-esm:1.263->3
|   |
|   |--(Q) : Sub=subscriber2@isp.net:ESM-SLA-PROF 2->1/2/lns-esm:1.263->2
|   |
|   |--(Q) : Sub=subscriber2@isp.net:ESM-SLA-PROF 2->1/2/lns-esm:1.263->1
|   |

|
===============================================================================
*A:LNS#
```

The advantage of using conventional H-QoS schedulers is that they can be applied universally on ingress and egress to provide a subscriber aggregate rate capability. The disadvantage of this approach is that the aggregate rate defined in the scheduler-policy (or overridden in the sub-profile) cannot be dynamically overridden from RADIUS using the QoS-override VSA (**Alc-Subscriber-QoS-Override**). If ingress H-QoS is not a requirement, but the ability to override the subscriber egress aggregate-rate is, then H-QoS should be implemented using an egress port-scheduler.

The egress port-scheduler is functionally the same as a conventional H-QoS scheduler in the manner with which it arbitrates bandwidth across its child queues. However, it has some notable differences:

- It is applied at the egress port level. Any queue that uses that egress port to which it is applied that is not explicitly mapped to a port-scheduler is considered an orphan queue. Orphan queues are not serviced by the port-scheduler until all of its child queues have been serviced.
- Unlike conventional H-QoS schedulers that include only Ethernet overhead, the port-scheduler includes Preamble and Inter-Frame Gap for each packet.
- It is supported only on Ethernet ports, and only on egress.
- The egress aggregate rate applied to the subscriber can be overridden from RADIUS.

The first bullet point above is significant from an LNS perspective. In general, after ESM handling, downstream traffic for subscribers egresses the system over a physical port. This is not the case for L2TP subscribers, which are passed through to the MS-ISA for L2TP encapsulation before egressing the LNS (and in fact could egress the system on any number of physical ports). It is therefore not possible to apply the port-scheduler policy to the egress port in the conventional manner, and what is needed is a mechanism to apply the port-scheduler policy to the logical internal ports that interface to the MS-ISA. To achieve this, an intermediate object known as a **port-policy** is used, which, when configured, references the **port-scheduler** policy, and which subsequently is applied to the relevant **lns-group**.

Create the port-scheduler-policy.

```
configure
    qos
        port-scheduler-policy "egress-port-scheduler" create
        exit
    exit
exit
```

Create the port-policy and reference the previously configured port-scheduler policy.

```
configure
    port-policy "isa-port-policy" create
        egress-scheduler-policy "egress-port-scheduler"
    exit
exit
```

Attach the port-policy to the lns-group containing the MS-ISA.

```
configure
    isa
        lns-group 1 create
            shutdown
            port-policy "isa-port-policy"
            no shutdown
        exit
    exit
exit
```

Once the **port-scheduler** policy and **port-policy** are in place, the subscriber QoS can reference it. The QoS configuration previously used for conventional H-QoS schedulers differs in both the **sap-egress** policy and **sub-profile** when an egress **port-scheduler** is used. The queues within the **sap-egress** policy are each configured to be parented to the egress port-scheduler using the **port-parent** keyword (as opposed the **parent** keyword used for conventional H-QoS schedulers).

```
configure
    qos
        sap-egress 10 create
            queue 1 create
                port-parent level 2 weight 20
            exit
            queue 2 best-effort create
                port-parent level 2 weight 80
            exit
            queue 3 expedite create
                port-parent cir-level 3
                rate 1024 cir 1024
            exit
            fc af create
                queue 2
            exit
            fc be create
                queue 1
            exit
```

```
            fc ef create
                queue 3
            exit
            dscp be fc "be"
            dscp ef fc "ef"
            dscp af31 fc "af"
        exit
    exit
exit
```

The sub-profile contains no reference to scheduler policies, but instead contains a per-subscriber egress aggregate rate in Kb/s, defined through the **agg-rate-limit** parameter.

```
configure
    subscriber-mgmt
        sub-profile "ESM-SUB-PROF"
            egress
                no scheduler-policy
                agg-rate-limit 8000
            exit
        exit
    exit
exit
```

Once again, the queues assigned to the subscriber through the above SAP-ingress/egress QoS policies, together with accumulative statistics can be viewed using the **show service active-subscribers subscriber <***name***> detail** command (real time rates can be seen using the **monitor** command). The scheduler SAP-egress queues mapped as child queues to a port-scheduler can be validated using the **show qos scheduler-hierarchy subscriber <***name***> egress** command. The **detail** argument as an extension of this command provides a significant amount of detail on bandwidth allocated to each queue by the scheduler in the within-CIR and above-CIR passes. It also provides a useful snapshot on offered traffic loads in Kb/s on a per-queue basis. Alternatively, all of the child queues and orphans mapped to the port-scheduler can be displayed using the **show qos scheduler-hierarchy port <***slot/mda/lns-esm***>** command, again with the optional **detail** argument.

```
*A:LNS# show qos scheduler-hierarchy subscriber "subscriber2@isp.net" egress

===============================================================================
Scheduler Hierarchy - Subscriber subscriber2@isp.net
===============================================================================
Egress Scheduler Policy :
-------------------------------------------------------------------------------
Root (Egr)
| slot(1)
|--(Q) : Sub=subscriber2@isp.net:ESM-SLA-PROF 2->1/2/lns-esm:1.263->3  (Port 1/2/lns
-esm)
|
|--(Q) : Sub=subscriber2@isp.net:ESM-SLA-PROF 2->1/2/lns-esm:1.263->2  (Port 1/2/lns
-esm)
|
```

```
|--(Q) : Sub=subscriber2@isp.net:ESM-SLA-PROF 2->1/2/lns-esm:1.263->1  (Port 1/2/lns
-esm)
|


===============================================================================
*A:LNS#
```

With the previously configured QoS policies and schedulers in place, the aggregate rate limit in use for the subscriber can be viewed using the **show service active-subscribers subscriber** <*name*> **detail** command. There are three fields in this output that are of interest here. The **E. Agg Rate Limit** field shows the configured rate-limit in the sub-profile and is therefore relatively static. The **RADIUS Rate-Limit** field shows the aggregate rate received by RADIUS using the **Alc-Subscriber-QoS-Override** VSA, which overrides any rate-limit statically configured in the sub-profile. Finally, the Oper-Rate-Limit shows the static or RADIUS-received rate-limit, minus any other H-QoS adjustments, such as Multicast H-QoS adjustment (snooping on IGMP joins) or ANCP line-rate adjustments.

```
*A:LNS# show service active-
subscribers subscriber "subscriber2@isp.net" detail | match expression " E. Agg Rate
 Limit|RADIUS Rate-Limit|Oper-Rate-Limit"
E. Sched. Policy : N/A                            E. Agg Rate Limit: 8000
RADIUS Rate-Limit: N/A
Oper-Rate-Limit  : 8000
*A:LNS#
```

Overriding the **agg-rate-limit** defined in the sub-profile can be done as part of the RADIUS Access-Accept, or through a Change of Authorization (CoA), and as previously outlined uses the **Alc-Subscriber-QoS-Override** VSA. This override function can be used, for example, to reconcile the LNS aggregate rate with the subscriber downstream rate learned through the TxConnectSpeed AVP in the ICCN message from the LAC. This ensures that the LNS does not overwhelm any downstream access node, and ensures that the LNS is responsible for all QoS scheduling in the event of congestion. In the following example, an override of the aggregate rate to 10Mb/s is sent as a CoA.

```
107 2016/05/27 08:20:32.84 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Change of Authorization(43) id 37 len 66 from 172.16.1.1:36767 vrid 1
    SESSION ID [44] 22 EA4CFF0000000A5748028E
    VSA [26] 16 Alcatel(6527)
      SUBSCRIBER QOS OVERRIDE [126] 14 e:r:rate=10000
"
```

Re-issuing the **show service active-subscribers subscriber** <*name*> **detail** command after the CoA shows that the **RADIUS Rate-Limit** field and the **Oper-Rate-Limit** field both correctly show 10Mb/s.

```
*A:LNS# show service active-
subscribers subscriber "subscriber2@isp.net" detail | match expression " E. Agg Rate
```

```
 Limit|RADIUS Rate-Limit|Oper-Rate-Limit"
E. Sched. Policy : N/A                              E. Agg Rate Limit: 8000
RADIUS Rate-Limit: 10000
Oper-Rate-Limit  : 10000
*A:LNS#
```

Propagating QoS Markings to L2TP/MPLS Headers

It is often desirable to mark the L2TP header (DCSP) or MPLS header (EXP) based on the class of service that is carried in the encapsulated subscriber IP payload. In general, when a packet is classified and mapped to an FC on ingress, that FC value is carried in the internal switch fabric header and is present when the packet is egressing the node. In the case of L2TP traffic however, the operation becomes a little more complex due to the fact that traffic transits the MS-ISA board with subscriber QoS implemented ‘mid-chassis’.

In the upstream direction (from subscriber to LNS), traffic arrives encapsulated in L2TP at the ingress IOM, and is passed through the MS-ISA via the internal object Ins-net. When traffic exits the MS-ISA through Ins-esm as native IP, it is subject to subscriber ingress QoS implemented on the (MS-ISA) carrier IOM. Traffic is classified and mapped to an FC at this point, and that FC mapping is maintained in the switch fabric header. As a result, marking of traffic is effected by the network egress QoS policy.

Figure 62 **Ingress/Egress QoS Processing**

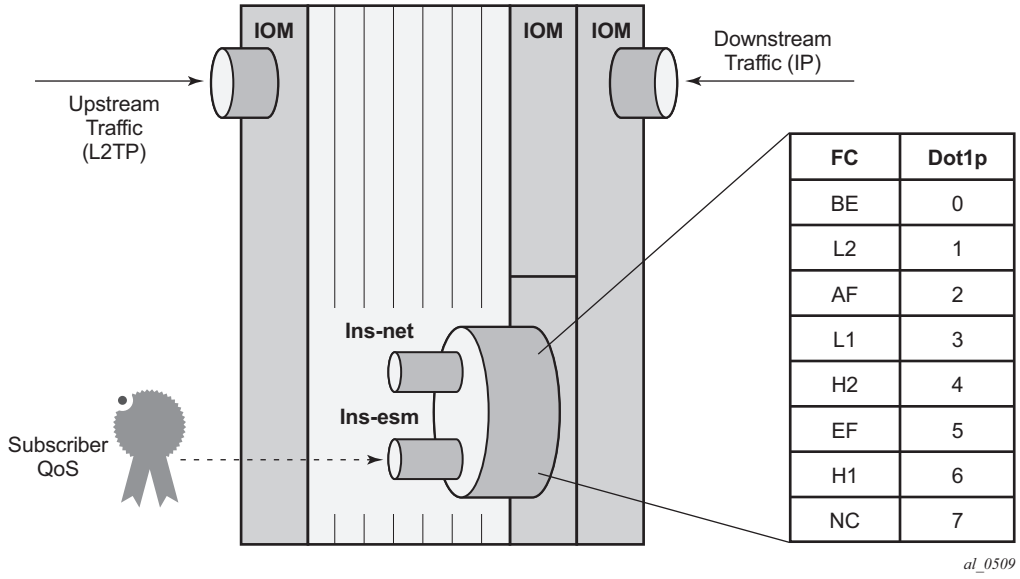

In the downstream direction (from LNS to subscriber), traffic arrives at the ingress IOM as native IP and is diverted to the MS-ISA via the lns-esm internal object. At the lns-esm, the traffic is subject to subscriber egress QoS. When the traffic is passed through lns-esm to the MS-ISA for L2TP encapsulation, internal Q-in-Q VLAN tags are attached as previously described. As the lns-esm is effectively a SAP-egress, the internal switch fabric header containing the FC marking is removed at this point, and, as a result FC information is lost.

Therefore, in order to allow for FC-continuity through the MS-ISA, the system implements a queue-group at the ingress of lns-net that has a dot1p to FC mapping as shown in the preceding figure. Assuming a SAP-egress QoS policy that employs FCs BE, AF and EF, the QoS policy would include the additional configuration to implement the appropriate dot1p marking as shown in the following output. When traffic arrives at lns-net, it is classified and mapped into the appropriate FCs, and the associated FC mapping included in the switch fabric header. At network egress, the L2TP packet is then subject to marking as defined in the network egress QoS policy.

```
configure
    qos
        sap-egress 10 create
            queue 1 create
            exit
            queue 2 best-effort create
            exit
            queue 3 expedite create
            exit
            fc af create
                queue 2
                dot1p 2
            exit
            fc be create
                queue 1
                dot1p 0
            exit
            fc ef create
                queue 3
                dot1p 5
            exit
            dscp be fc "be"
            dscp ef fc "ef"
            dscp af31 fc "af"
        exit
    exit
exit
```

Framed-Route

The majority of residential services in broadband networks have a single registered 32-bit IPv4 address on the WAN side of the RG and a private (RFC 1918) network on the LAN side. Traffic from the LAN toward the BNG (and Internet) is thereafter subject to Network Address and Port Translation (NAPT). However, a common requirement for delivery of business services is the ability for the BNG to recognize one or more IP subnets on the LAN side of the RG that is not subject to NAT, and the subscriber prefix is a route to a network. This is achieved using the standard RADIUS **Framed-Route** attribute, or dynamic BGP peering. Both serve the function of allowing one or more subnets to be learned at the LNS with a next-hop IP address of the RG WAN.

To provide an example of the use of Framed-Route, the Retail VRF VPRN 2 is again used, and in fact requires no modification in order to support subscribers with Framed-Routes. In general ESM, where Framed-Route is used, there is a requirement to configure **anti-spoof type nh-mac**, but for LNS SAPs this is the default. The RADIUS users file is updated to also return a Framed-Route attribute for prefix 10.128.46.0/24 with a next-hop determined by the subscriber IP prefix. The prefix has a metric of 10, and has a tag of value 200, which may be used for example, for routing policy.

```
subscriber2@isp.net     Cleartext-Password := "letmein"
                        Alc-Subsc-ID-Str = "subscriber2@isp.net",
                        Alc-Subsc-Prof-Str = "ESM-SUB-PROF",
                        Alc-SLA-Prof-Str = "ESM-SLA-PROF",
                        Alc-Serv-Id = "2",
                        Alc-Interface = "VPRN2-GROUP-INT",
                        Service-Type = Framed-User,
                        Framed-Protocol = PPP,
                        Framed-IP-Address = 10.10.148.22,
                        Framed-Route = "10.128.46.0/24 0.0.0.0 10 tag 200",
```

In SR OS, a prefix learned through the Framed-Route attribute is known internally as a **Managed Route**. Once the subscriber is instantiated, the presence of the Managed Route can be verified as installed.

```
*A:LNS# show service id 2 ppp session detail | match "Managed Routes" post-lines 5
Managed Routes
-------------------------------------------------------------------------------
IP Address                                Status       Metric Tag        Pref
-------------------------------------------------------------------------------
10.128.46.0/24                            installed    10     200        0
-------------------------------------------------------------------------------
*A:LNS#
```

The Managed Route can also be seen present in the VPRN routing-table, learned through protocol **Managed**.

```
*A:LNS# show router 2 route-table protocol managed

===============================================================================
Route Table (Service: 2)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age         Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.128.46.0/24                                  Remote  Managed   00h00m32s  0
       10.10.148.22                                                  10
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
*A:LNS#
```

L2TP Tunnel Switching (LTS)

In general, L2TP tunnels are established directly between LAC and LNS. However, if there are a large number of LAC devices (and therefore a large number of L2TP tunnels), it may be desirable to perform some aggregation of these tunnels before presenting them to the LNS. This is implemented by one or more LNS devices performing the function of an L2TP Tunnel Switch (LTS). The LTS terminates multiple L2TP tunnels from the LAC(s), and sources a single L2TP tunnel toward the target LNS, switching L2TP sessions from one tunnel to another tunnel accordingly.

SR OS supports LTS functionality, and from a configuration perspective, it requires no more than L2TP being placed in a no shutdown state in the required routing context if the relevant attributes are returned from RADIUS. As with LNS functions, at least one MS-ISA is required to support LTS functions. In this example, VPRN 1 is used with the previously defined configuration. To recap, this VPRN has a single L2TP group “L2TP-GROUP-1”, and within that group, a single tunnel defined “L2TP-TUNNEL-1” that terminates the tunnel from the LAC. To demonstrate LTS functionality, the LAC continues to function as a LAC, the LNS functions as an LTS, and PE-1 becomes the LNS.

The RADIUS users file for subscriber1@isp.net is modified to include a number of additional attributes and VSAs. As previously described, the **Alc-Serv-Id** and **Alc-Interface** define the service ID and group-interface where the subscriber is terminated, and this can be any IES or VPRN service. The **Alc-Tunnel-Serv-Id** VSA is used to identify the service from which the L2TP tunnel is initiated. This can be the same service in which the subscriber is terminated, or it can be a different service. If it is a different service, then the minimum requirement is that L2TP is placed in a no shutdown state. In this example, the subscriber is terminated in service VPRN 1 and the L2TP tunnel is also initiated from service VPRN 1. The remaining attributes are

standard attributes defined in RFC 2868 for L2TP tunnel set-up. The Tunnel-Assignment-Id attribute is used to maintain the concept of groups and tunnels, where Tunnel-Assignment-Id:0 is used to indicate the group name and Tunnel-Assignment-Id:1 is used to indicate the tunnel name. This provides sufficient information for the LTS to be able to initiate an L2TP tunnel without any requirement for nodal configuration above that already configured.

```
subscriber1@isp.net     Cleartext-Password := "letmein"
                        Alc-Subsc-ID-Str = "subscriber1@isp.net",
                        Alc-Subsc-Prof-Str = "ESM-SLA-PROF",
                        Alc-SLA-Prof-Str = "ESM-SUB-PROF",
                        Alc-Serv-Id = "1",
                        Alc-Interface = "LNS-GROUP-INT",
                        Alc-Tunnel-Serv-Id = 1,
                        Tunnel-Assignment-Id:0 = "RADIUS-returned-Tunnel-Group",
                        Tunnel-Type:1 += L2TP,
                        Tunnel-Medium-Type:1 += IP,
                        Tunnel-Server-Endpoint:1 += 192.168.0.3,
                        Tunnel-Password:1 += "password",
                        Tunnel-Assignment-Id:1 += "RADIUS-returned-Tunnel-Name",
                        Tunnel-Client-Auth-Id = "LTS",
```

The LAC forwards the PPP session into the LAC to LTS tunnel, and after the LTS receives the ICCN from the LAC, it proceeds in authenticating the subscriber. RADIUS returns the above attributes with sufficient information for the LTS to instantiate the subscriber and initiates an L2TP tunnel/session with PE-1, the target LNS. The LNS then authenticates the user once more, this time providing it with IP address information through IPCP negotiation. This interaction between PPP client and LNS is transparent to the LTS, which is responsible for switching PPP packets between L2TP sessions. However, the user is instantiated in the system as a fully-fledged subscriber.

```
*A:LNS# show service id 1 ppp session

===============================================================================
PPP sessions for service 1
===============================================================================
User-Name
  Descr.
          Up Time       Type  Termination     IP/L2TP-Id/Interface-Id MC-Stdby
-------------------------------------------------------------------------------
subscriber1@isp.net
  vprn:1 connid:281630942 tid:4297 sid:22750
          0d 00:00:30   oL2tp lac             32660054
-------------------------------------------------------------------------------
No. of PPP sessions: 1
===============================================================================
*A:LNS#
```

Within VPRN 1, there are two L2TP tunnels active. The first entry with Connection Id 32636928 belongs to group **RADIUS-returned-Tunnel-Group** (derived from RADIUS attribute Tunnel-Assignment-Id:0) and has tunnel name **RADIUS-returned-Tunnel-Name** (derived from RADIUS attribute Tunnel-Assignment-Id:1). This is the tunnel from LTS to LNS, and it is in the **Established** state and has one session active. The second entry with Connection Id 281608192 is the statically defined tunnel from the LAC, belonging to the CLI-configured group L2TP-GROUP-1 with tunnel name L2TP-TUNNEL-1. This tunnel is also in the **Established** state, with one session active.

```
*A:LNS# show router 1 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State              Blacklist-state   Ses Active
  Group                                                             Ses Total
    Assignment
-------------------------------------------------------------------------------
32636928   498       12390     established        not-blacklisted   1
  RADIUS-returned-Tunnel-
Group                                      1
    RADIUS-returned-Tunnel-Name
281608192  4297      11865     established        not-blacklisted   1
  L2TP-GROUP-
1                                                      2
    L2TP-TUNNEL-1
-------------------------------------------------------------------------------
No. of tunnels: 2
===============================================================================
*A:LNS#
```

Equally, within VPRN 1, there are two L2TP sessions active for subscriber subscriber1@isp.net. Session ID 32660054 is carried in Tunnel-ID 498, which, as shown in the previous output, is the tunnel toward the LNS, while session ID 281630942 is carried in Tunnel-ID 4297, which is the tunnel toward the LAC.

```
*A:LNS# show router 1 l2tp session

===============================================================================
L2TP Session Summary
===============================================================================
ID                  Control Conn ID     Tunnel-ID   Session-ID  State
-------------------------------------------------------------------------------
32660054            32636928            498         23126       established
281630942           281608192           4297        22750       established
  subscriber1@isp.net
  interface: LNS-GROUP-INT
  service-id: 1
  32660054
-------------------------------------------------------------------------------
No. of sessions: 2
===============================================================================
*A:LNS#
```

IPv6

The deployment of IPv6 into residential broadband networks dictates some design choices, or perhaps even some enforced IPv6 address allocation mechanisms. Bridged or Routed Residential Gateways (RGs). Numbered or unnumbered WAN. Stateful (DHCPv6) or stateless (Stateless Address Auto-Configuration, or SLAAC) address assignment. The purpose of this example is not to show every possibility, but simply to demonstrate that enabling IPv6 is possible at the LNS, just as if this were a conventional BNG doing PPP Termination and Aggregation (PTA). This example uses a widely adopted approach of dual-stack Routed RG with DHCPv6 Prefix Delegation.

The configuration of VPRN 2 is modified to include some IPv6 parameters. In the subscriber-interface the **delegated-prefix-len** parameter is set to **variable** to indicate that prefixes delegated to subscribers may be of varying length (the default delegated prefix length is /64). The **allow-unmatching-prefixes** parameter tells the subscriber-interface to operate in an IPv6 unnumbered mode, allowing IPv6 addresses to be allocated to subscribers that do not fall within the range of any IPv6 subnet defined under the subscriber-interface. Within the group-interface, the **ipv6** context places router-advertisements into a no shutdown state and has the **managed-configuration** flag set to indicate that stateful (DHCPv6) address configuration is to be used.

There is also a **dhcp6 proxy-server** enabled, that provides an interworking function between RADIUS (where the Delegated Prefix will be learned from) and the DHCPv6 client. The proxy will take the RADIUS-provided prefix and responds to the clients Solicit message with an DHCPv6 Advertise message containing the delegated prefix (IA_PD). Because the DHCPv6 messages from the client need to be received over the subscriber PPP session, the proxy-server is configured to allow this using the **client-applications ppp** command. Finally, there is a static-route to black-hole the /48 IPv6 prefix. The client is allocated a /64 prefix from this range and this static-route is used to provide an aggregated upstream prefix advertisement.

```
configure
    service
        vprn 2
            subscriber-interface "VPRN2-SUB-INT" create
                ipv6
                    default-dns 2001:db8:2c41::56
                    delegated-prefix-len variable
                    allow-unmatching-prefixes
                exit
                group-interface "VPRN2-GROUP-INT" lns create
                    ipv6
                        router-advertisements
                            managed-configuration
                            no shutdown
                        exit
                        dhcp6
```

```
                            proxy-server
                                client-applications ppp
                                no shutdown
                            exit
                        exit
                    exit
                exit
            exit
            static-route-entry 2a00:8010:1b00::/48
                black-hole
                    no shutdown
                exit
            exit
            no shutdown
        exit
    exit
exit
```

The RADIUS users file entry for subscriber2@isp.net is also modified to return the IPv6 Delegated Prefix using the standard attribute **Delegated-IPv6-Prefix**.

```
subscriber2@isp.net     Cleartext-Password := "letmein"
                        Alc-Subsc-ID-Str = "subscriber2@isp.net",
                        Alc-Subsc-Prof-Str = "ESM-SUB-PROF",
                        Alc-SLA-Prof-Str = "ESM-SLA-PROF",
                        Alc-Serv-Id = "2",
                        Alc-Interface = "VPRN2-GROUP-INT",
                        Service-Type = Framed-User,
                        Framed-Protocol = PPP,
                        Framed-IP-Address = 10.10.148.22,
                        Delegated-IPv6-Prefix = 2001:db8:1b00:100::/64
```

After the PPP LCP phase and RADIUS authentication, the LNS is aware that the subscriber also has IPv6 enabled (in this case because it received the **Delegated-IPv6-Prefix** attribute). As a result, the LNS begins to negotiate both IPCP and IPv6CP with the client. For IPv6CP, only an Interface-ID is negotiated, for which the LNS uses an EUI-64 extended version of the chassis MAC address. Once IPv6CP negotiation is completed, the client can initiate a DHCPv6 Solicit for a delegated prefix (IA_PD option). After a successful Advertise/Request/Reply exchange the subscriber is instantiated as dual-stack IPv4/IPv6.

```
*A:LNS# show service active-subscribers subscriber "subscriber2@isp.net"

===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber subscriber2@isp.net (ESM-SUB-PROF)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[1/2/lns-esm:1.263] - sla:ESM-SLA-PROF
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
```

```
10.10.148.22
                00:00:10:c9:5b:0c    PPP 1          IPCP         2          Y
2001:db8:1b00:100::/64
                00:00:10:c9:5b:0c    PPP 1          DHCP6-PD     2          Y
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
*A:LNS#
```

Conclusion

SR OS offers a comprehensive feature set for LNS implementations. The MS-ISA provides the hardware-assist for L2TP encapsulation/de-capsulation while the carrier IOM implements conventional subscriber management functions.

Multi-Chassis IPSec Redundancy

This chapter provides information about multi-chassis IPSec redundancy configurations.

Topics in this chapter include:

Applicability

This chapter is applicable to 7750 SR-7/12/12e with IOM3-XP or IMMs and chassis mode D and the 7450 ESS-6/7/12 with IOM3-XP or IMM in mixed mode. See the release notes for a complete overview of supported hardware.

This chapter was originally written for and tested on release 10.0.R8. The CLI in this version corresponds to 13.0.R6.

Overview

Multi-Chassis IPSec redundancy (MC-IPSec) is a stateful inter-chassis IPSec failover mechanism. IPSec tunnel states are synchronized between the master and standby chassis. A tunnel-group failure on the master or a master chassis failure could trigger MC-IPSec failover to the standby chassis.

The following are some highlights of this feature:

- Internet Key Exchange version 2 (IKEv2) only
- Multi-active tunnel-group only
- The granularity of failover is tunnel-group, which means a specific tunnel-group could failover to the standby chassis independent of other tunnel-groups on the master chassis
- Supports both static and dynamic LAN-to-LAN tunnel

This feature has the following building blocks:

- Master election
  - MIMP (MC-IPSec Mastership Protocol) runs between the chassis to elect a master, MIMP run for each tunnel-group independently
- Synchronization
  - MCS (Multi-Chassis Synchronization) synchronizes IPSec states between chassis
- Routing
  - MC-IPSec-aware routing attracts traffic to the master chassis
  - Shunting support
  - MC-IPSec aware Virtual Router Redundancy Protocol (VRRP)

Figure 63 **MC-IPSec Architecture**

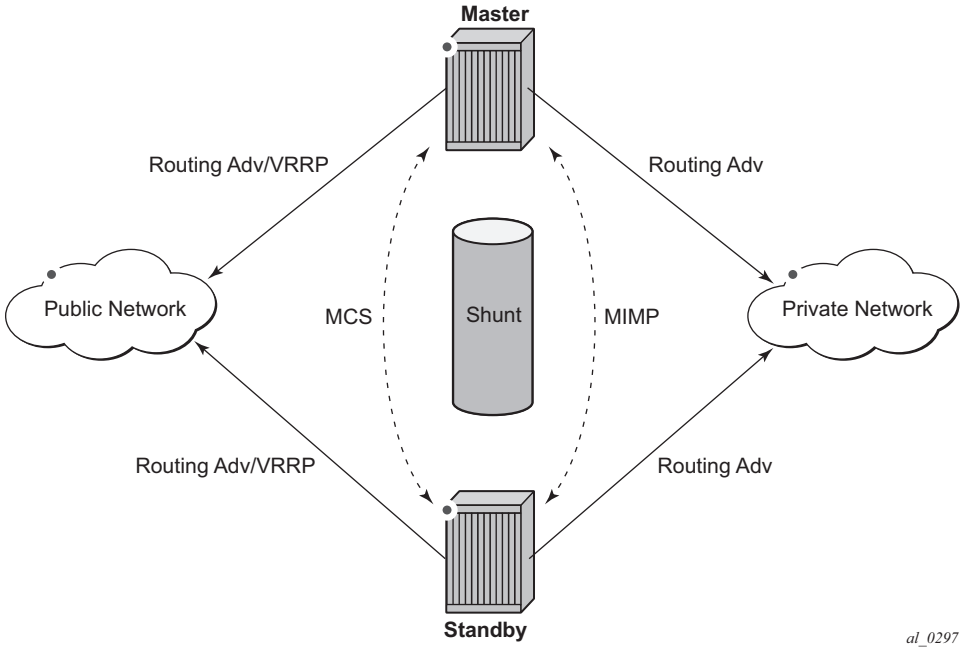

The fundamentals of MC-IPSec are:

- Only the master processes Encapsulating Security Payload (ESP) and IKE traffic. If the standby receives traffic, it could shunt it to the master if possible. The traffic will be discarded if the standby fails to shunt the traffic.
- The same local gateway address should be provisioned on both chassis.

- MC-IPSec does not synchronize configurations.
- MC-IPSec aware routing attracts traffic to the master for both public and private services. This is achieved by exporting the corresponding IPSec routes to the routing protocol via a route policy and setting a different routing metric according to the MC-IPSec state.
- In case of a Layer 2 public network, MC-IPSec aware VRRP could be used to trigger VRRP switchover upon MC-IPSec switchover.
- MCS synchronizes IPSec states between chassis so that existing IPSec tunnels do not need to be re-established upon switchover.
- MIMP elects mastership between two chassis, and it could also detect chassis failure and tunnel-group failure; a central BFD session could be bound to the MIMP to achieve fast chassis failure detection.

Configuration

The test topology is shown in Figure 64.

Figure 64 **Test Topology**

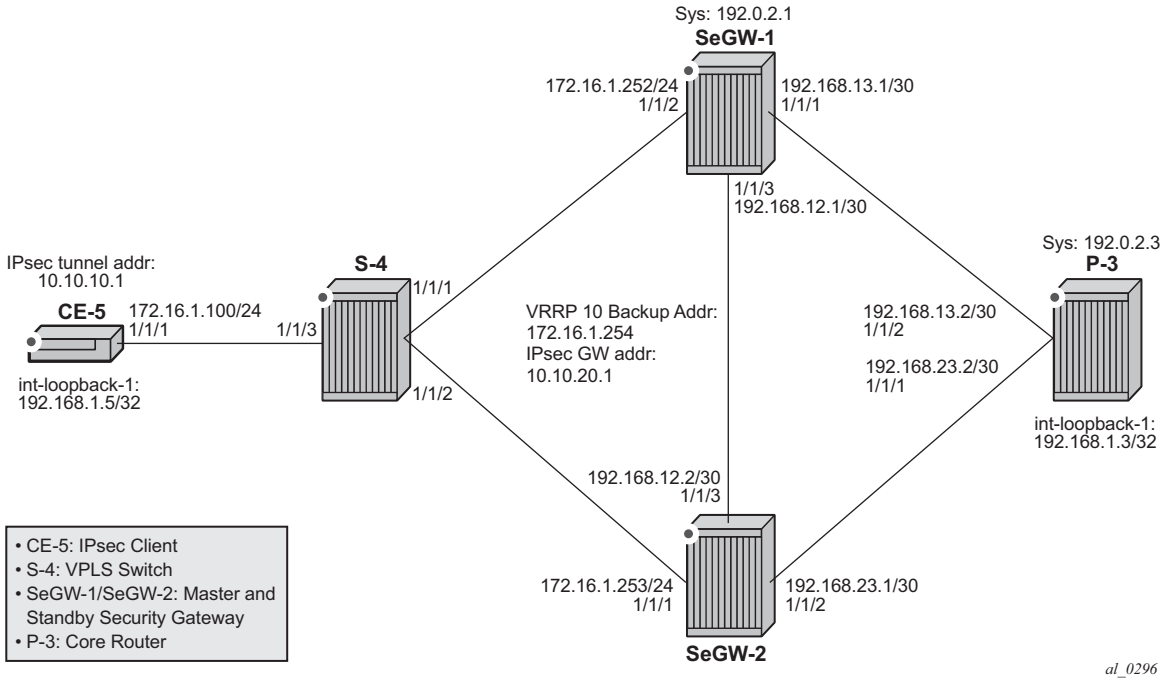

Test setup:

- An IPSec tunnel is initiated by CE-5 and terminated on the master of SeGW-1/SeGW-2.
- IES 1 and VPRN 2 are the public and private services, respectively, on SeGW-1/SeGW-2/CE-5.
- VPRN 2 is also configured on P-3.
- Static LAN-to-LAN tunnel with pre-shared key.
- Local VPLS service 3 on S-4 to simulate a Layer 2 switch.
- VRRP 10 between SeGW-1 and SeGW-2 to provide a backup address 192.168.1.254, which is also the default next-hop for CE-5.
- VRRP policy 1 is bound to VRRP 10 to change the in-use priority upon MC-IPSec switchover.
- OSPF is the IGP running in the base routing instance between SeGW-1, SeGW-2 and P-3.
- MP-BGP is running between SeGW-1, SeGW-2 and P-3 for exchanging VPRN 2 routes.
- A ping between loopback interface address: 192.168.1.5 on CE-5 and 192.168.1.3 on P-3 in VPRN 2 is used to verify the connectivity over the IPSec tunnel.

The MC-IPSec configuration commands are shown below.

```
config>redundancy>multi-chassis>
     peer <ip-address> [create]
          sync
              ipsec
              tunnel-group <tunnel-group-id> sync-tag <tag-name> [create]
          mc-ipsec
              bfd-enable
              discovery-interval <interval-1> [boot <interval-2>]
              hold-on-neighbor-failure <multiplier>
              keep-alive-interval <interval>
              tunnel-group <tunnel-group-id> [create]
                  peer-group <tunnel-group-id>
                  priority <priority>
                  shutdown

config>router>policy-options>policy-statement>entry>from>
     state ipsec-master-with-peer|ipsec-non-master|ipsec-master-without-peer
     protocol ipsec

config>service>ies>if>
config>service>vprn>if>
     static-tunnel-redundant-next-hop <ip-address>
     dynamic-tunnel-redundant-next-hop <ip-address>

config>isa>tunnel-grp>
     ipsec-responder-only

config>vrrp>policy>priority-event>
     mc-ipsec-non-forwarding <tunnel-grp-id>
```

```
hold-clear <seconds>
hold-set <seconds>
priority <priority-level> explicit
```

Parameters:

- **peer <ip-address> [create] —** This command creates or enters a multi-chassis peer. The peer address is by default the system address. This can be changed on the peer using the config>redundancy>multi-chassis>peer>source-address command.
- **sync**>**ipsec** — This command enables MCS to synchronize IPSec states.
- **tunnel-group** <*tunnel-group-id*> **sync-tag** <*tag-name*> [**create**] — This command enables MCS to synchronize the IPSec states of the specified tunnel-group.The **sync-tag** parameter is used to match the tunnel-group of the peer. The tunnel-group states with the same sync-tag on both chassis will be synchronized.
- **mc-ipsec** — This command enters the multi-chassis IPSec configuration context.
- **bfd-enable** — This command enables tracking a central BFD session, if the BFD session goes down, then the system considers the peer is down and changes the mc-ipsec status of the configured tunnel-group accordingly.

  The BFD session uses the source address of MCS as its source address and the MCS peer address as the destination address. Other BFD parameters are configured with the **bfd** command on the interface that the MCS source address resides on.

  Configuration of this command is optional.
- **discovery-interval** <*interval-1*> [**boot** <*interval-2*>] — This command specifies the time interval that the tunnel-group stays in “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 discovery-interval after system boot-up, it is optional, and when it is not specified, the value for interval-1will be used. Both intervals have a default value of 300 seconds.
- **hold-on-neighbor-failure** <*multiplier*> — This command specifies the number of keep-alive failures before considering the peer to be down. Default is 3.
- **keep-alive-interval** <*interval*> — This command specifies the time interval of the mastership election protocol keep-alive packets. Default value is 1 second, range: 0.5 ~ 50 seconds.
- **tunnel-group** <*tunnel-group-id*> [**create**] — This command enables multi-chassis redundancy for the specified tunnel-group, or enters an already configured tunnel-group context. The configured tunnel-groups could failover independently.

- **peer-group** <*tunnel-group-id*> — This command specifies the corresponding tunnel-group id on the peer node. The peer tunnel-group id is not necessarily equal to local tunnel-group id.
- **priority** <*priority*> — This command specifies the local priority of the tunnel-group, this is used to elect a master, where the higher number wins. If the priorities are the same, then the peer which has more active ISAs wins; if the priority and the number of active ISAs are same, then the peer with higher IP address wins. Default value is 100, range: 0..255
- **shutdown** — This command disables the multi-chassis redundancy for the specified tunnel-group
- **state ipsec-master-with-peer|ipsec-non-master|ipsec-master-without-peer** — These commands specify the mc-ipsec state in a “from” statement of a route policy entry.

  ipsec-master-with-peer: The corresponding tunnel-group is Master with peer reachable.

  ipsec-master-without-peer: The corresponding tunnel-group is Master with peer unreachable.

  ipsec-non-master: The corresponding tunnel-group is **not** Master.
- **protocol ipsec** — This command specifies the IPSec as protocol in a “from” statement of a route policy entry. **protocol ipsec** means the /32 local gateway routes (of both static and dynamic tunnels) and reverse route of dynamic tunnel.
- **static-tunnel-redundant-next-hop** <*ip-address*>

  dynamic-tunnel-redundant-next-hop <*ip-address*> — These commands specify the redundant next-hop address on a public or private IPSec interface (with public or private tunnel-sap) for a static and dynamic IPSec tunnel respectively. The specified next-hop address will be used by the standby node to shunt traffic to the master in case it receives any traffic.

  The next-hop address will be resolved in the routing table of the corresponding service.

  Notes:

  - Shunting is supported over:
    - Directly connected SAP
    - Spoke SDP terminated IP interface
  - Shunting over auto-bind tunnel is not supported.
  - Shunting will not work if the tunnel-group is down.
- **ipsec-responder-only** — With this command configured, the system will only act as IKE responder except for the automatic CHILD_SA rekey upon MC-IPSec switchover.

  This command is required for MC-IPSec support of static LAN-to-LAN tunnel

- **mc-ipsec-non-forwarding** <*tunnel-grp-id*> — This command creates a new VRRP policy priority event: **mc-ipsec-non-forwarding**. It will be triggered whenever the specified tunnel-group enters non-forwarding state.
- **hold-clear** <*seconds*> — This command configures hold time before clearing the event. Default value is 0 seconds. Range: 0..86400 seconds
- **hold-set** <*seconds*> — This command configures hold time before setting the event. Default value is 0 seconds. Range: 0..86400 seconds
- **priority** <*priority-level*> **explicit** — This command sets the VRRP in-use priority to the configured value upon the event. Default value is 0, range: 0..254

Before starting

- The system time of SeGW-1 and SeGW-2 must be the same. Otherwise, this feature will not work. Using a time sync protocol like NTP/SNTP is the recommended method.
- SeGW-1 and SeGW-2 must be IP reachable in the base routing instance because both MCS and MIMP run in the base routing instance.

Step 0: Configure CE-5.

- IES 1 and VPRN 2 are the public and private service.
- A static default route points to the VRRP backup address 172.16.1.254.
- A static IPSec tunnel “tunnel-1” has local address 10.10.10.1 and remote address 10.10.20.1.
- A loopback interface in VPRN 2 with address 192.168.1.5/32 is used as source address for the ping traffic in later step.
  - The ping traffic is used to test the connectivity between CE-5 and P-3 over IPSec tunnel “tunnel-1”.

```
*A:CE-5# configure router
        interface "int-CE-5-S-4"
            address 172.16.1.100/24
            port 1/1/1
        exit
        interface "system"
        exit
        autonomous-system 64496
        static-route 0.0.0.0/0 next-hop 172.16.1.254
        exit

*A:CE-5# configure ipsec
        ike-policy 1 create
            ike-version 2
            dpd
        exit
        ipsec-transform 1 create
```

```
        exit

*A:CE-5# configure isa
        tunnel-group 1 create
            primary 1/2
            no shutdown
        exit

*A:CE-5# configure service
        ies 1 customer 1 create
            interface "int-IPsec-Public-1" create
                address 10.10.10.254/24
                tos-marking-state untrusted
                sap tunnel-1.public:1 create
                exit
            exit
            no shutdown
        exit


*A:CE-5# configure service
        vprn 2 customer 1 create
            ipsec                security-policy 1 create
                    entry 10 create
                        local-ip 192.168.1.5/32
                        remote-ip 192.168.1.3/32
                    exit
                exit
            exit
            route-distinguisher 64496:2
            interface "int-loopback-1" create
                address 192.168.1.5/32
                loopback
            exit
            interface "int-IPsec-private-1" tunnel create
                sap tunnel-1.private:1 create
                    ipsec-tunnel "tunnel-1" create
                        security-policy 1
                        local-gateway-address 10.10.10.1 peer 10.10.20.1
                                              delivery-service 1
                        dynamic-keying
                            ike-policy 1
                            pre-shared-key "ALU"
                            transform 1
                        exit
                        no shutdown
                    exit
                exit
            exit
            static-route 192.168.1.3/32 ipsec-tunnel "tunnel-1"
            no shutdown
        exit
```

Step 1. Configure S-4.

- A local VPLS service 3 simulates a Layer 2 switch between CE-5, SeGW-1 and SeGW-2.

```
*A:S-4# configure service
       vpls 3 customer 1 create
           sap 1/1/1 create
           exit
           sap 1/1/2 create
           exit
           sap 1/1/3 create
           exit
           no shutdown
       exit
```

Step 2. Configure P-3

- P-3 simulates the core network router, which connects to both SeGW-1 and SeGW-2.
- A loopback interface with address 192.168.1.3/32 in VPRN 2 is the destination address of the ping traffic from CE-5.
- MP-BGP session between P-3 and SeGW-1/SeGW-2 to receive 192.168.1.5/32 route in VPRN 2.
- GRE spoke SDPs to connect to SeGW-1 and SeGW-2.

```
*A:P-3# configure router
       interface "int-P-3-SeGW-1"
           address 192.168.13.2/30
           port 1/1/2
       exit
       interface "int-P-3-SeGW-2"
           address 192.168.23.2/30
           port 1/1/1
       exit
       interface "system"
           address 192.0.2.3/32
       exit
       autonomous-system 64496
       exit

*A:P-3# configure router
       ospf
           area 0.0.0.0
               interface system
               exit
               interface "int-P-3-SeGW-1"
               exit
               interface "int-P-3-SeGW-2"
               exit
           exit
       exit

*A:P-3# configure service
```

```
        sdp 31 create
            signaling off
            far-end 192.0.2.1
            no shutdown
        exit
        sdp 32 create
            signaling off
            far-end 192.0.2.2
            no shutdown
        exit

*A:P-3# configure service
        vprn 2 customer 1 create
            route-distinguisher 64496:2
            vrf-target target:64496:2
            interface "int-loopback-1" create
                address 192.168.1.3/32
                loopback
            exit
            spoke-sdp 31 create
                description "SDP to SeGW-1"
            exit
            spoke-sdp 32 create
                description "SDP to SeGW-2"
            exit
            no shutdown
        exit

*A:P-3# configure router
        bgp
            group "MPBGP"
                family vpn-ipv4
                type internal
                neighbor 192.0.2.1
                exit
                neighbor 192.0.2.2
                exit
            exit
            no shutdown
        exit
```

Step 3. Configure IPSec tunnel on SeGW-1.

- The tunnel-group must be in multi-active mode before MC-IPSec can be enabled.
- Interface "int-Redundant-1" is a spoke-sdp terminated interface is used for shunting.
- GRE SDPs 12 and 13 toward SeGW-2 and P-3.
- IPSec tunnel “tunnel-1” is the tunnel to CE-5; both SeGW-1 and SeGW-2 use the same local gateway address: 10.10.20.1.

```
*A:SeGW-1# configure isa
        tunnel-group 1 create
```

```
            ipsec-responder-only
            multi-active
            mda 1/2
            no shutdown
        exit

*A:SeGW-1# configure router
        interface "int-SeGW-1-P-3"
            address 192.168.13.1/30
            port 1/1/1
        exit
        interface "int-SeGW-1-SeGW-2"
            address 192.168.12.1/30
            port 1/1/3
        exit
        interface "system"
            address 192.0.2.1/32
            bfd 100 receive 100 multiplier 3
        exit
        autonomous-system 64496
        static-route 10.10.10.0/24 next-hop 172.16.1.100
        exit

*A:SeGW-1# configure router
        ospf
            area 0.0.0.0
                interface "system"
                exit
                interface "int-SeGW-1-SeGW-2"
                exit
                interface "int-SeGW-1-P-3"
                exit
            exit
        exit

*A:SeGW-1# configure ipsec
        ike-policy 1 create
            ike-version 2
            ipsec-lifetime 7200
            isakmp-lifetime 172800
        exit
        ipsec-transform 1 create
        exit


*A:SeGW-1# configure service
        sdp 12 create
            signaling off
            far-end 192.0.2.2
            no shutdown
        exit
        sdp 13 create
            signaling off
            far-end 192.0.2.3
            no shutdown
        exit

*A:SeGW-1# configure service
        ies 1 customer 1 create
```

```
            interface "int-SeGW-1-S-4" create
                address 172.16.1.252/24
                sap 1/1/2 create
                exit
            exit
            interface "int-IPsec-Public-1" create
                address 10.10.20.254/24
                tos-marking-state untrusted
                sap tunnel-1.public:1 create
                exit
                static-tunnel-redundant-next-hop 192.168.12.2
            exit
            no shutdown
        exit

*A:SeGW-1# configure service
        vprn 2 customer 1 create
            ipsec
                security-policy 1 create
                    entry 10 create
                        local-ip 192.168.1.3/32
                        remote-ip 192.168.1.5/32
                    exit
                exit
            exit
            route-distinguisher 64496:2
            vrf-target target:64496:2
            interface "int-IPsec-Private-1" tunnel create
                sap tunnel-1.private:1 create
                    ipsec-tunnel "tunnel-1" create
                        security-policy 1
                        local-gateway-address 10.10.20.1 peer 10.10.10.1
                                              delivery-service 1
                        dynamic-keying
                            ike-policy 1
                            pre-shared-key "ALU"
                            transform 1
                        exit
                        no shutdown
                    exit
                exit
                static-tunnel-redundant-next-hop 192.168.120.2
            exit
            interface "int-Redundant-1" create
                address 192.168.120.1/30
                spoke-sdp 12:20 create
                    ingress
                        vc-label 2049
                    exit
                    egress
                        vc-label 2048
                    exit
                    no shutdown
                exit
            exit
            static-route 192.168.1.5/32 ipsec-tunnel "tunnel-1"
            spoke-sdp 12 create
                description "SDP to SeGW-2"
            exit
```

```
            spoke-sdp 13 create
                description "SDP to P-3"
            exit
            no shutdown
        exit
```

Step 4. Enable MC-IPSec for tunnel-group 1 on SeGW-1

- Create a multi-chassis peer using the system address of SeGW-2.
- Enable MCS for IPSec and tunnel-group 1.
- Enable MC-IPSec for the tunnel-group with a configured priority 200.
- Bind a central BFD session to MC-IPSec from the system interface.

```
*A:SeGW-1# configure redundancy
        multi-chassis
            peer 192.0.2.2 create
                sync
                    ipsec
                    tunnel-group 1 sync-tag "tunnel-group-1" create
                    no shutdown
                exit
                mc-ipsec
                    bfd-enable
                    tunnel-group 1 create
                        peer-group 1
                        priority 200
                        no shutdown
                    exit
                exit
                no shutdown
            exit
        exit

*A:SeGW-1# configure router
        interface "system"
            address 192.0.2.1/32
            bfd 100 receive 100 multiplier 3
            no shutdown
        exit
```

Step 5. Configure MC-IPSec aware routing on SeGW-1.

- Export static route 192.168.1.5/32 in VPRN 2 to P-3 by using route-policy "IPsec-to-MPBGP".
- Set the local preference of the 192.168.1.5/32 according to the MC-IPSec state:
  - ipsec-master-with-peer: 200
  - ipsec-non-master:100
  - ipsec-master-without-peer: 200

State “ipsec-master-without-peer” could be used to attract traffic to the designated master in case of “Dual Master” (meaning two chassis lose the MIMP connection in the base routing instance). In this example, SeGW-1 has local preference 200 and SeGW-2 has local preference 100 for ipsec-master-without-peer.

– Apply the policy “IPsec-to-MPBGP” in VPRN 2.

```
*A:SeGW-1# configure router
        policy-options
            begin
            prefix-list "CE-5-Internal"
                prefix 192.168.1.5/32 exact
            exit
            community "vprn2" members "target:64496:2"
            policy-statement "IPsec-to-MPBGP"
                entry 10
                    from
                        prefix-list "CE-5-Internal"
                        state ipsec-master-with-peer
                    exit
                    action accept
                        community add "vprn2"
                        local-preference 200
                    exit
                exit
                entry 20
                    from
                        prefix-list "CE-5-Internal"
                        state ipsec-non-master
                    exit
                    action accept
                        community add "vprn2"
                        local-preference 100
                    exit
                exit
                entry 30
                    from
                        prefix-list "CE-5-Internal"
                        state ipsec-master-without-peer
                    exit
                    action accept
                        community add "vprn2"
                        local-preference 200
                    exit
                exit
                default-action accept
                    community add "vprn2"
                exit
            exit
            commit
        exit
    exit


*A:SeGW-1# configure router
        bgp
            group "MPBGP"
```

```
                family vpn-ipv4
                type internal
                neighbor 192.0.2.2
                exit
                neighbor 192.0.2.3
                exit
            exit
            no shutdown
        exit

*A:SeGW-1# configure service
        vprn 2 customer 1 create
            vrf-export "IPsec-to-MPBGP"
        exit
```

Step 6. Configure MC-IPSec-aware VRRP on SeGW-1.

- The VRRP instance needs to be in preempt mode.
- Use “mc-ipsec-non-forwarding” priority event to lower the in-use VRRP priority upon MC-IPSec switchover, which ensures VRRP and MC-IPSec have the same master.
- Apply the vrrp-policy on interface "int-SeGW1-S1" of IES 1.
  - This only needs to be configured on the designated VRRP master, in this case, SeGW-1.

```
*A:SeGW-1# configure vrrp
        policy 1
            priority-event
                mc-ipsec-non-forwarding 1
                    priority 50 explicit
                exit
            exit
        exit


*A:SeGW-1# configure service
        ies 1 customer 1 create
            interface "int-SeGW-1-S-4" create
                vrrp 10
                    backup 172.16.1.254
                    priority 200
                    policy 1
                    ping-reply
                exit
                sap 1/1/2 create
                exit
            exit
```

Step 7. Repeat Step 3 to Step 5 on SeGW-2.

```
*A:SeGW-2# configure isa
        tunnel-group 1 create
            ipsec-responder-only
            multi-active
```

```
            mda 1/2
            no shutdown
        exit

*A:SeGW-2# configure redundancy
        multi-chassis
            peer 192.0.2.1 create
                sync
                    ipsec
                    tunnel-group 1 sync-tag "tunnel-group-1" create
                    no shutdown
                exit
                mc-ipsec
                    bfd-enable
                    tunnel-group 1 create
                        peer-group 1
                        priority 150
                        no shutdown
                    exit
                exit
                no shutdown
            exit
        exit

*A:SeGW-2# configure router
        interface "int-SeGW-2-P-3"
            address 192.168.23.1/30
            port 1/1/2
        exit
        interface "int-SeGW-2-SeGW-1"
            address 192.168.12.2/30
            port 1/1/3
        exit
        interface "system"
            address 192.0.2.2/32
            bfd 100 receive 100 multiplier 3
        exit
        autonomous-system 64496
        static-route 10.10.10.0/24 next-hop 172.16.1.100
        exit

*A:SeGW-2# configure router
        ospf
            area 0.0.0.0
                interface "system"
                exit
                interface "int-SeGW-2-SeGW-1"
                exit
                interface "int-SeGW-2-P-3"
                exit
            exit
        exit

*A:SeGW-2# configure ipsec
        ike-policy 1 create
            ike-version 2
            ipsec-lifetime 7200
            isakmp-lifetime 172800
        exit
```

```
        ipsec-transform 1 create
        exit

*A:SeGW-2# configure router
        policy-options
            begin
            prefix-list "CE-5-Internal"
                prefix 192.168.1.5/32 exact
            exit
            community "vprn2" members "target:64496:2"
            policy-statement "IPsec-to-MPBGP"
                entry 10
                    from
                        prefix-list "CE-5-Internal"
                        state ipsec-master-with-peer
                    exit
                    action accept
                        community add "vprn2"
                        local-preference 200
                    exit
                exit
                entry 20
                    from
                        prefix-list "CE-5-Internal"
                        state ipsec-non-master
                    exit
                    action accept
                        community add "vprn2"
                        local-preference 100
                    exit
                exit
                entry 30
                    from
                        prefix-list "CE-5-Internal"
                        state ipsec-master-without-peer
                    exit
                    action accept
                        community add "vprn2"
                        local-preference 100
                    exit
                exit
                default-action accept
                    community add "vprn2"
                exit
            exit
            commit
        exit

*A:SeGW-2# configure router
        bgp
            group "MPBGP"
                family vpn-ipv4
                type internal
                neighbor 192.0.2.1
                exit
                neighbor 192.0.2.3
                exit
            exit
            no shutdown
```

```
        exit

*A:SeGW-2# configure service
        sdp 21 create
            signaling off
            far-end 192.0.2.1
            no shutdown
        exit
        sdp 23 create
            signaling off
            far-end 192.0.2.3
            no shutdown
        exit

*A:SeGW-2# configure service
        ies 1 customer 1 create
            interface "int-SeGW-2-S-4" create
                address 172.16.1.253/24
                vrrp 10
                    backup 172.16.1.254
                    ping-reply
                exit
                sap 1/1/1 create
                exit
            exit
            interface "int-IPsec-Public-1" create
                address 10.10.20.254/24
                tos-marking-state untrusted
                sap tunnel-1.public:1 create
                exit
                static-tunnel-redundant-next-hop 192.168.12.1
            exit
            no shutdown
        exit

*A:SeGW-2# configure service
        vprn 2 customer 1 create
            ipsec
                security-policy 1 create
                    entry 10 create
                        local-ip 192.168.1.3/32
                        remote-ip 192.168.1.5/32
                    exit
                exit
            exit
            vrf-export "IPsec-to-MPBGP"
            route-distinguisher 64496:2
            vrf-target target:64496:2
            interface "int-IPsec-Private-1" tunnel create
                sap tunnel-1.private:1 create
                    ipsec-tunnel "tunnel-1" create
                        security-policy 1
                        local-gateway-address 10.10.20.1 peer 10.10.10.1
                                              delivery-service 1
                        dynamic-keying
                            ike-policy 1
                            pre-shared-key "ALU"
                            transform 1
                        exit
```

```
                    no shutdown
                exit
            exit
            static-tunnel-redundant-next-hop 192.168.120.1
        exit
        interface "int-Redundant-1" create
            address 192.168.120.2/30
            spoke-sdp 21:20 create
                ingress
                    vc-label 2048
                exit
                egress
                    vc-label 2049
                exit
                no shutdown
            exit
        exit
        static-route 192.168.1.5/32 ipsec-tunnel "tunnel-1"
        spoke-sdp 21 create
            description "SDP to SeGW-1"
        exit
        spoke-sdp 23 create
            description "SDP to P-3"
        exit
        no shutdown
    exit
```

Step 8. Verify the MC-IPSec status on SeGW-1 and SeGW-2.

- Verify that SeGW-1 is the master and SeGW-2 is the standby for tunnel-group 1 because SeGW-1 has higher priority 200.
- Verify that SeGW-1 is the VRRP 10 master and SeGW-2 is the backup.

```
*A:SeGW-1# show redundancy multi-chassis mc-ipsec peer 192.0.2.2

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.2
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 11/25/2015 09:04:17

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             master
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
*A:SeGW-1#
```

```
*A:SeGW-2# show redundancy multi-chassis mc-ipsec peer 192.0.2.1

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.1
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 11/25/2015 10:04:50

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              150       Up             standby
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
*A:SeGW-2#



*A:SeGW-1# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-1-S-4                   10    No  Up   Master       200        1
                                 IPv4      Up   1           200        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
*A:SeGW-1#


*A:SeGW-2# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-2-S-4                   10    No  Up   Backup       100        1
                                 IPv4      Up   n/a         100        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
*A:SeGW-2#
```

Step 9. Trigger the tunnel-1 setup on CE-5 by sending pings.

```
*A:CE-5# ping router 2 192.168.1.3
PING 192.168.1.3 56 data bytes
64 bytes from 192.168.1.3: icmp_seq=2 ttl=63 time=1.95ms.
64 bytes from 192.168.1.3: icmp_seq=3 ttl=63 time=1.88ms.
64 bytes from 192.168.1.3: icmp_seq=4 ttl=63 time=1.89ms.
64 bytes from 192.168.1.3: icmp_seq=5 ttl=63 time=1.90ms.
Request timed out. icmp_seq=1.

---- 192.168.1.3 PING Statistics ----
5 packets transmitted, 4 packets received, 20.00% packet loss
round-trip min = 1.88ms, avg = 1.91ms, max = 1.95ms, stddev = 0.026ms
*A:CE-5#


*A:CE-5# show ipsec tunnel

===============================================================================
IPsec Tunnels
===============================================================================
TunnelName                       LocalAddress      SvcId        Admn   Keying
  SapId                            RemoteAddress     DlvrySvcId   Oper   Sec
                                                                         Plcy
-------------------------------------------------------------------------------
tunnel-1                         10.10.10.1        2            Up     Dynamic
  tunnel-1.private:1               10.10.20.1        1            Up     1
-------------------------------------------------------------------------------
IPsec Tunnels: 1
===============================================================================
*A:CE-5#
```

Step 10. Verify that the tunnel status on SeGW-1/SeGW-2 is “up”.

- Verify that MCS database is in-sync, so the tunnel status is “up” on both chassis.
- Verify P-3 receives two 192.168.15/32 VPN IPv4 routes, the route from SeGW-1 has local preference 200, and the one from SeGW-2 has 100.

```
*A:SeGW-1# show ipsec tunnel
===============================================================================
IPsec Tunnels
===============================================================================
TunnelName                       LocalAddress      SvcId        Admn   Keying
  SapId                            RemoteAddress     DlvrySvcId   Oper   Sec
                                                                         Plcy
-------------------------------------------------------------------------------
tunnel-1                         10.10.20.1        2            Up     Dynamic
  tunnel-1.private:1               10.10.10.1        1            Up     1
-------------------------------------------------------------------------------
IPsec Tunnels: 1
===============================================================================
*A:SeGW-1#
```

```
*A:SeGW-2# show ipsec tunnel
===============================================================================
IPsec Tunnels
===============================================================================
TunnelName                       LocalAddress      SvcId        Admn   Keying
  SapId                            RemoteAddress     DlvrySvcId   Oper   Sec
                                                                         Plcy
-------------------------------------------------------------------------------
tunnel-1                         10.10.20.1        2            Up     Dynamic
  tunnel-1.private:1               10.10.10.1        1            Up     1
-------------------------------------------------------------------------------
IPsec Tunnels: 1
===============================================================================
*A:SeGW-2#


*A:SeGW-2# show redundancy multi-chassis sync
===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 192.0.2.1
Description             : (Not Specified)
Authentication          : Disabled
Source IP Address       : 192.0.2.2
Admin State             : Enabled
Warm standby            : No
Remote warm standby     : No
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     : IPsec
Sync Admin State        : Up
Sync Oper State         : Up
Sync Oper Flags         :
DB Sync State           : inSync
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
===============================================================================
===============================================================================
*A:SeGW-2#


*A:P-3# show router bgp routes vpn-ipv4
===============================================================================
 BGP Router ID:192.0.2.3        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
```

```
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop (Router)                                 Path-Id     Label
      As-Path
-------------------------------------------------------------------------------
u*>i  64496:2:192.168.1.5/32                           200         None
      192.0.2.1                                        None        262143
      No As-Path
*i    64496:2:192.168.1.5/32                           100         None
      192.0.2.2                                        None        262143
      No As-Path
u*>i  64496:2:192.168.120.0/30                         100         None
      192.0.2.1                                        None        262143
      No As-Path
*>i   64496:2:192.168.120.0/30                         100         None
      192.0.2.2                                        None        262143
      No As-Path
-------------------------------------------------------------------------------
Routes : 4
===============================================================================
*A:P-3#
```

Step 11. Trigger MC-IPSec switchover by shutting down the MS-ISA.

- Verify the VRRP/MC-IPSec state on SeGW-1 is “master”, SeGW-2 is “backup”/”standby”.
- Shutdown the MS-ISA on SeGW-1, which is currently Master.
- Verify that the MC-IPSec state of tunnel-group 1 on SeGW-1 becomes “notEligible”, SeGW-2 becomes “master”.

  Note: notEligible means the tunnel-group is down, refer to the SR OS MS-ISA Guide for details description of MIMP states.
- Verify that the VRRP state on SeGW-1 becomes “backup” and SeGW-2 becomes “master”. This is triggered by MC-IPSec switchover, configured via mc-ipsec-non-forwarding event in vrrp-policy 1.

```
*A:SeGW-1# show redundancy multi-chassis mc-ipsec peer 192.0.2.2

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.2
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 11/25/2015 13:48:18

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
```

```
ID              Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------
1               1              200       Up             master
-------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
=========================================================================
===============================================================================
*A:SeGW-1#


*A:SeGW-1# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-1-S-4                   10    No  Up   Master       200       1
                                 IPv4      Up   1           200        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
*A:SeGW-1#


*A:SeGW-2# show redundancy multi-chassis mc-ipsec peer 192.0.2.1

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.1
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 11/25/2015 13:49:49

=========================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
=========================================================================
ID              Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------
1               1              150       Up             standby
-------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
=========================================================================
===============================================================================
*A:SeGW-2#


*A:SeGW-2# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
```

```
-------------------------------------------------------------------------------
int-SeGW-2-S-4                  10    No  Up   Backup       100       1
                                IPv4      Up   n/a          100       No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
*A:SeGW-2#


*A:SeGW-1# configure card 1 mda 2 shutdown


*A:SeGW-1# show redundancy multi-chassis mc-ipsec peer 192.0.2.2

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.2
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 11/25/2015 13:48:18

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             notEligible
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
*A:SeGW-1#



*A:SeGW-1# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                  VR Id Own Adm  State       Base Pri   Msg Int
                                IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-1-S-4                  10    No  Up   Backup       200       1
                                IPv4      Up   1           50         No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
*A:SeGW-1#



*A:SeGW-2# show redundancy multi-chassis mc-ipsec peer 192.0.2.1
```

```
===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.1
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 11/25/2015 13:49:49

===========================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===========================================================================
ID             Peer Group     Priority  Admin State    Mastership
---------------------------------------------------------------------------
1              1              150       Up             master
---------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===========================================================================
===============================================================================
*A:SeGW-2#
```

```
*A:SeGW-2# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                    VR Id Own Adm  State       Base Pri   Msg Int
                                  IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-2-S-4                    10    No  Up   Master       100       1
                                  IPv4      Up   n/a         100       No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
*A:SeGW-2#
```

Step 12. Trigger the MC-IPSec switchover by rebooting SeGW-1.

- Restore state as in Step 10 (before the MC-IPSec switchover).
  - Note: The MC-IPSec switchover could be triggered manually with the **tools perform redundancy multi-chassis mc-ipsec force-switchover tunnel-group 1** command.
- Verify the VRRP/MC-IPSec state on SeGW-1 is “master”, SeGW-2 is “backup”/”standby”.
- Reboot SeGW-1 which is the current Master.
- Verify the MC-IPSec state of tunnel-group 1 on SeGW-2 becomes “eligible” while SeGW-1 is rebooting.

– Verify the VRRP state on SeGW-2 becomes “master” during SeGW-1 reboot.
– After SeGW-1 comes up, verify MC-IPSec state of tunnel-group 1 is “discovery” initially, and then becomes “standby”;
  • Note: The “discovery” state means system has not established the MIMP session with peer yet.
– Verify the MC-IPSec state of tunnel-group 1 on SeGW-2 becomes “master” when SeGW-1 becomes “standby”.
– After SeGW-1 comes up, verify the VRRP state is “backup”.

```
A:SeGW-1# tools perform redundancy multi-chassis mc-ipsec force-switchover tunnel-
group 1
Forcing a mastership switchover may impact traffic. Are you sure (y/n)? y
A:SeGW-1#

A:SeGW-1# show redundancy multi-chassis mc-ipsec peer 192.0.2.2

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.2
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 11/26/2015 09:56:11

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             master
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
A:SeGW-1#


A:SeGW-1# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                    VR Id Own Adm  State       Base Pri   Msg Int
                                  IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-1-S-4                    10    No  Up   Master       200       1
                                  IPv4      Up   1           200        No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
```

```
A:SeGW-1#


*A:SeGW-2# show redundancy multi-chassis mc-ipsec peer 192.0.2.1

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.1
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 11/25/2015 13:49:49

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              150       Up             standby
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
*A:SeGW-2#


*A:SeGW-2# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-2-S-4                   10    No  Up   Backup       100       1
                                 IPv4      Up   n/a         100       No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
*A:SeGW-2#

A:SeGW-1# admin reboot
Are you sure you want to reboot (y/n)? y

*A:SeGW-2# show redundancy multi-chassis mc-ipsec peer 192.0.2.1

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.1
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 11/25/2015 13:49:49
```

```
===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              150       Up             eligible
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
*A:SeGW-2#

*A:SeGW-2# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-2-S-4                   10    No  Up   Master       100       1
                                 IPv4      Up   n/a         100       No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
*A:SeGW-2#
```

Then SeGW-1 comes up.

```
A:SeGW-1# show redundancy multi-chassis mc-ipsec peer 192.0.2.2

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.2
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 11/26/2015 10:38:44

===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Admin State    Mastership
-------------------------------------------------------------------------------
1              1              200       Up             discovery
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
A:SeGW-1#


A:SeGW-1# show redundancy multi-chassis mc-ipsec peer 192.0.2.2
```

```
===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.2
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 11/26/2015 10:38:44

===========================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===========================================================================
ID             Peer Group     Priority  Admin State    Mastership
---------------------------------------------------------------------------
1              1              200       Up             standby
---------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===========================================================================
===============================================================================
A:SeGW-1#


A:SeGW-1# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-1-S-4                   10    No  Up   Backup       200        1
                                 IPv4      Up   1           50         No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
A:SeGW-1#

*A:SeGW-2# show redundancy multi-chassis mc-ipsec peer 192.0.2.1

===============================================================================
Multi-Chassis MC-IPsec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 192.0.2.1
Keep Alive Intvl: 1.0 secs            Hold on Nbr Fail     : 3
Discovery Intvl : 300 secs            Discovery Boot Intvl : 300 secs
BFD             : Enable
Last update     : 11/25/2015 13:49:49

===========================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===========================================================================
ID             Peer Group     Priority  Admin State    Mastership
---------------------------------------------------------------------------
1              1              150       Up             master
---------------------------------------------------------------------------
```

```
Multi Active Tunnel Group Entries found: 1
===============================================================================
===============================================================================
*A:SeGW-2#

*A:SeGW-2# show router vrrp instance

===============================================================================
VRRP Instances
===============================================================================
Interface Name                   VR Id Own Adm  State       Base Pri   Msg Int
                                 IP        Opr  Pol Id      InUse Pri  Inh Int
-------------------------------------------------------------------------------
int-SeGW-2-S-4                   10    No  Up   Master       100       1
                                 IPv4      Up   n/a         100       No
  Backup Addr: 172.16.1.254
-------------------------------------------------------------------------------
Instances : 1
===============================================================================
*A:SeGW-2#
```

Configuration Guidelines

The following is a list of configuration and operational guidelines that the user should follow for MC-IPSec:

- To avoid high CPU load and issues in some complex cases, the following are suggestions for configuring IKEv2 lifetime:
  1. Both IKE_SA and CHILD_SA lifetime on MC-IPSec chassis (SeGW-1 and SeGW-2) should be around 3 times larger than on the IPSec peer (CE-5).
  2. With the first rule, the lifetime of the side with smaller lifetime should NOT be too small (these being the default values):
     - IKE_SA: >= 86400 seconds
     - CHILD_SA: >= 3600 seconds
  3. With the first rule, on the side with smaller lifetime, the IKE_SA lifetime should be at least 3 times larger than CHILD_SA lifetime.
- IKE protocol is the control plane of IPSec, so IKE packet should be treated as high QoS priority in end-to-end path of public service.
  - On public interface, a sap-ingress qos policy should be configured to ensure IKE packet gets high QoS priority.
- Configure responder-only under tunnel-group for static LAN-to-LAN tunnel.
- Enable DPD (Dead Peer Detection) on peer side, configure “no dpd” on MC-IPSec chassis side.

- Direct and redundant physical link between MC-IPSec chassis should be configured with enough bandwidth for MCS and shunting traffic, and proper QoS configuration to make sure the MIMP/MCS packet treated as high priority traffic.
- System time must be same on both MC-IPSec chassis.
- Check and make sure the protection status is "nominal" on both chassis before you do a controlled switchover. Protection status could be displayed via command “show redundancy multi-chassis mc-ipsec peer <addr>”.
- Wait at least 5 minutes between two consecutive switchovers if possible, to prevent a second switchover happening before the standby is ready to take over mastership.

Conclusion

MC-IPSec provides a stateful multi-chassis IPSec redundancy solution. This is very important in a carrier grade network, especially in applications like mobile backhaul where high value 3G/4G mobile services run over IPSec tunnels.

NAT in Combination with ESM

This chapter provides information about Network Address Translation (NAT) in combination with Enhanced Subscriber Management (ESM).

Topics in this chapter include:

Applicability

This chapter is applicable to SROS routers and was tested on release 8.0.R.6 and retested on 14.0.R1. See the release notes for a complete list of supported hardware. Chassis mode B or higher is required.

SROS supports Source Network Address and Port Translation (SNAPT aka N:1) and Source Network Address Translation (SNAT aka 1:1) to provide continuity of legacy IPv4 services during the migration to native IPv6.

The router can operate in two different NAT-modes known as:

- Large Scale NAT
- Layer 2-aware NAT

Layer 2 aware is also known as NAT in combination with Enhanced Subscriber Management (ESM), and is the only NAT mode covered in this chapter.

Overview

Layer 2-Aware NAT performs source address and port translation as commonly deployed for shared Internet access. NAT in SR OS is used to provide consumer broadband or business Internet customers access to IPv4 Internet resources with a shared pool of IPv4 addresses, such as may occur with the IPv4 exhaustion.

TCP/UDP connections use ports for multiplexing, with 65536 ports available for every IP address. When many hosts are trying to share a single public IP address there is a chance of port collision where two different hosts use the same source port for a connection. The resulting collision is avoided in SNAPT devices by translating the source port and tracking this in a stateful manner. All SNAPT devices are stateful in nature and must monitor connection establishment and traffic to maintain translation mappings.

In most circumstances, SNAPT requires the inside host to establish a connection to the public Internet host or server before a mapping and translation will occur. With the initial outbound IP packet, SNAPT knows the inside IP, inside port, remote IP, remote port and protocol. L2-aware NAT will also take into account the subscriber identification string. With this information the SNAPT device can select an IP and port combination (referred to as outside IP and outside port) from its pool of addresses and create a unique mapping for this data flow.

Any traffic returned from the server will use the outside IP and outside port in the destination IP/port fields matching the unique NAT mapping. The mapping then provides the inside IP and inside port for translation.

Figure 65 **Network Address Translation Overview**

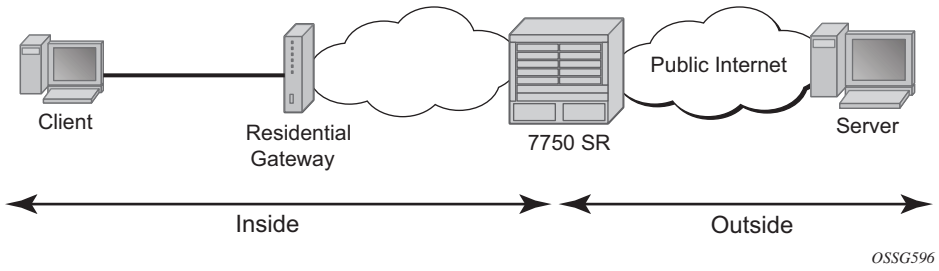

L2-aware NAT supports the following ESM hosts:

- IP over Ethernet (IPoE)
- PPP over Ethernet (PPPoE)
- L2TP Network Server (LNS)

L2-aware NAT is not supported on static- or arp-hosts.

L2-aware NAT differentiates between two interfaces; the inside interface, toward the residential gateway and, the outside interface, toward the public network, as seen in Figure 65. The outside IP needs to be a public IP address.

NAT is supported in the base and VPRN routing contexts. NAT traffic can enter in a VPRN routing context, and leave through a base or VPRN routing context. L2-aware NAT allows IP address re-use toward residential customers.

A typical flow-session will be recorded using the following fields:

- Subscriber identification string
- Inside IP- Outside IP
- Inside port- Outside port
- Inside VRFid- Outside VRFid

This chapter will focus on the NAT configuration and functionality. For completeness other configuration will be given, but not explained in detail. Two IPoE clients will be set up with the same IP address inside one VPRN. One PPPoE client will be set up using a VPRN different from the public VPRN.

Server Functionality Behind NAT

Servers (such as HTTP, SMTP, etc) or Peer-to-Peer (P2P) applications can have difficulty when operating behind SNAPT because traffic from the Internet can reach the NAT without a mapping present in the SNAPT.

Different methods can be employed to overcome this, including:

- Port forwarding
- Session Traversal Utilities for NAT (STUN) support
- Application Layer Gateways (ALG)

SROS supports all three methods following the best-practices RFC for TCP (RFC 5382, *NAT Behavioral Requirements for TCP*) and UDP (RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP*). Port forwarding, allows for servers using well-known ports lower than1024 (such as HTTP and SMTP) to request the appropriate outside port for permanent allocation.

STUN is facilitated by the support of Endpoint-Independent Filtering and Endpoint-Independent Mapping (RFC 4787) in the NAT device, allowing STUN-capable applications to detect the NAT and allow inbound P2P connections for that specific application. Many new voice over IP clients and instant messaging chat applications are STUN capable.

Application Layer Gateways (ALGs) allow the NAT to monitor the application running over TCP or UDP and make appropriate changes in the NAT translations accordingly. SR OS implements ALGs for FTP, SIP and RTSP.

Configuration

Hardware Configuration

L2-aware NAT is implemented using the MS-ISA MDA hosted by an IOM3-XP. The MS-ISA card is a multi-purpose MDA which can be used for multiple applications like LNS, video (FCC/RET/VQE), AA (application assurance/DPI), tunneling (GRE/IPSec), etc. This approach allows re-deploying the same hardware in a different software configuration for other purposes once the IPv4 to IPv6 transition is completed.

To support L2-aware NAT, the MS-ISA must be configured as an ISA-BB (Broadband) MDA. The ISA-BB can run multiple applications simultaneously. For example, L2-aware NAT can be combined with carrier grade-NAT.

```
configure
    card 1
        card-type iom3-xp
        --- snipped ---
        mda 2
            mda-type isa-bb
            no shutdown
        exit
        no shutdown
    exit
exit
```

An ISA NAT-group needs to be created. A NAT-group can host up to 6 MDAs for load-sharing or providing resilience when an MS-ISA fails. Multiple NAT groups can be created for example to achieve hardware segregation between residential and business customers.

```
configure isa
        nat-group 1 create
            description "L2 aware NAT group"
            active-mda-limit 1
            mda 1/2
            no shutdown
        exit
    exit
exit
```

The active-mda-limit controls the number of MS-ISA MDAs which can be used as active members of the NAT group. Each active card will be assigned sessions/flows and will process traffic. The backup cards are cold standby; they are used only in case of a failure of one (or more) of the active cards. Load balancing over the active cards is basedon the source IP address for the upstream direction, and on the outside destination IPaddress for the downstream direction. Public IP address pools are assigned to a dedicated card, thus resulting in both upstream and downstream traffic flowing through the same MS-ISA card.

All MS-ISA cards (active + backup) need to be configured under the ISA NAT group.

Service Configuration

Figure 66 **Setup Topology**

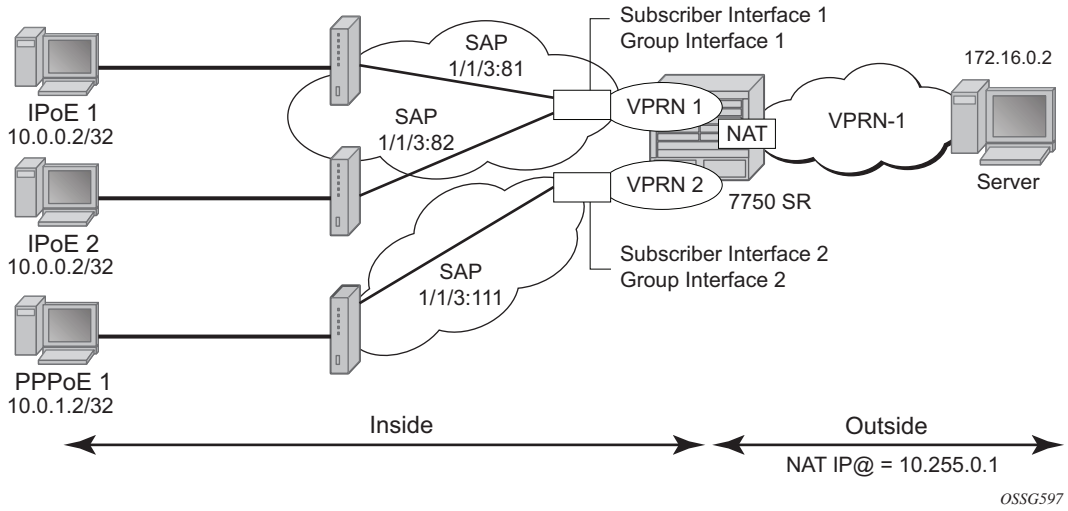

The example setup is shown in Figure 66 and Figure 67. There are three clients:

1. IPoE_1 going to SAP 1/1/3:81 in VPRN 1 (NAT to VPRN 1)
2. IPoE_2 going to SAP 1/1/3:82 in VPRN 1 (NAT to VPRN 1)
3. PPPoE_1 going to SAP 1/1/3:111 in VPRN 2 (NAT to VPRN 1)

Figure 67 **Simplified Routing Topology**

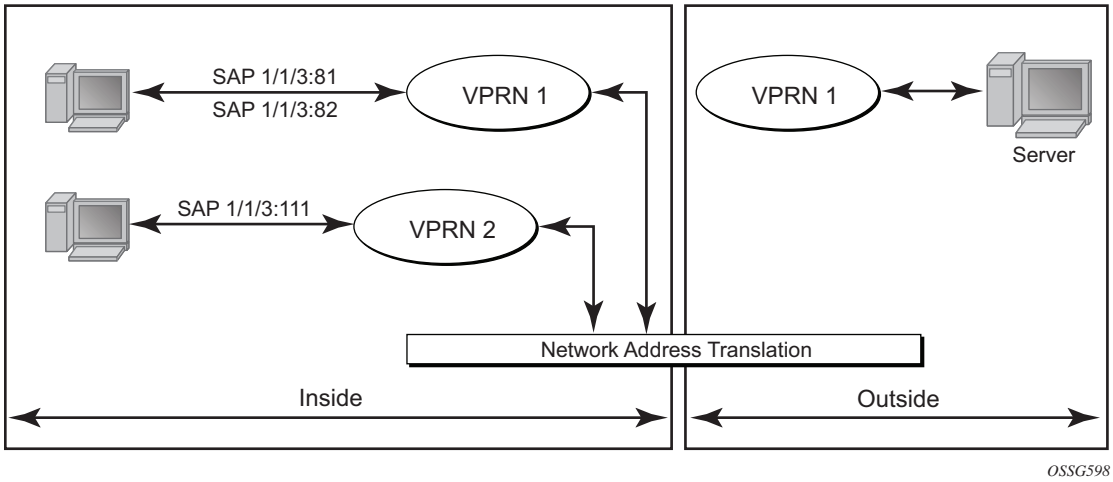

Initial Service and Enhanced Subscriber Management Configuration

The initial service and ESM configuration is shown for VPRN 1. This VPRN contains subscriber interface sub-int-1 with group interface group-int-ipoe-cpe, in routed central office (CO). There are two SAPs under the group interface, 1/1/1:81 and 1/1/1:82. The subscribers are IP over Ethernet subscribers. Upon receiving a DHCP request the subscriber will be authenticated using RADIUS. Because a proxy scenario is used, an emulated-server address is configured under the group interface.

```
configure
    service
        vprn 1 customer 1 create
            route-distinguisher 64496:1
            interface "int-PE-1-servers" create
                address 172.16.2.1/30
                sap 1/1/4:2 create
                exit
            exit
            interface "int-TEST" create
                address 10.11.11.1/32
                loopback
            exit
            subscriber-interface "sub-int-1" create
                address 10.0.0.254/24
                address 10.0.1.254/24
                group-interface "group-int-dhcp-cpe" create
                    mac 00:00:00:00:00:01
                    arp-populate
                    dhcp
```

```
                        proxy-server
                            emulated-server 10.0.0.254
                            lease-time hrs 1
                            no shutdown
                        exit
                        trusted
                        lease-populate 10
                        gi-address 10.0.0.254
                        no shutdown
                    exit
                    authentication-policy "radius-1"
                    sap 1/1/1:81 create
                        anti-spoof nh-mac
                        sub-sla-mgmt
                            def-sub-profile "sub-profile-nat"
                            def-sla-profile "sla-profile-nat"
                            sub-ident-policy "sub-ident-nat"
                            multi-sub-sap 10
                            no shutdown
                        exit
                    exit
                    sap 1/1/1:82 create
                        anti-spoof nh-mac
                        sub-sla-mgmt
                            def-sub-profile "sub-profile-nat"
                            def-sla-profile "sla-profile-nat"
                            sub-ident-policy "sub-ident-nat"
                            multi-sub-sap 10
                            no shutdown
                        exit
                    exit
                exit
            exit
            no shutdown
        exit
    exit
exit
```

All parameters are returned by the RADIUS server, including all ESM strings as well as the IP address, mask, and default gateway.

The RADIUS user configuration is as follows. The user's mac-address is used to authenticate the IPoE and PPPoE users.

```
00:0c:29:9d:10:2d    Cleartext-Password := "letmein"
                     Alc-Subsc-ID-Str = "ipoe-sub-00:0c:29:9d:10:2d",
                     Alc-SLA-Prof-Str = "sla-profile-nat",
                     Alc-Subsc-Prof-Str = "sub-profile-nat",
                     Framed-IP-Address = 10.0.0.2,
                     Framed-IP-Netmask = 255.255.255.0,
                     Alc-Default-Router = 10.0.0.254,

00:0c:29:34:cc:74    Cleartext-Password := "letmein"
                     Alc-Subsc-ID-Str = "ipoe-sub-00:0c:29:34:cc:74",
                     Alc-SLA-Prof-Str = "sla-profile-nat",
                     Alc-Subsc-Prof-Str = "sub-profile-nat",
                     Framed-IP-Address = 10.0.0.2,
```

```
                    Framed-IP-Netmask = 255.255.255.0,
                    Alc-Default-Router = 10.0.0.254,

00:0c:29:1d:44:34   Cleartext-Password := "letmein"
                    Alc-Subsc-ID-Str = "pppoe-sub-00:0c:29:1d:44:34",
                    Alc-SLA-Prof-Str = "sla-profile-nat",
                    Alc-Subsc-Prof-Str = "sub-profile-nat",
                    Framed-IP-Address = 10.0.1.2,
                    Framed-IP-Netmask = 255.255.255.0,
                    Alc-Default-Router = 10.0.0.254,
```

The subscriber management policies are as follows.

```
configure
    subscriber-mgmt
        sla-profile "sla-profile-nat" create
        exit
        sub-profile "sub-profile-nat" create
            nat-policy "nat-l2aware-vprn1"
            radius-accounting
                policy "nat-accounting"
            exit
        exit
        sub-ident-policy "sub-ident-nat" create
            sub-profile-map
                use-direct-map-as-default
            exit
            sla-profile-map
                use-direct-map-as-default
            exit
            app-profile-map
                use-direct-map-as-default
            exit
        exit
        ppp-policy "pppoe-nat" create
            max-sessions-per-mac 10
            ppp-authentication pap
        exit
    exit
exit
```

The RADIUS authentication and accounting policies are defined as follows:

```
configure
    router
        radius-server
            server "radius-172.16.1.2" address 172.16.1.2 secret vsecret1 create
                accept-coa
            exit
        exit
    exit
exit

configure
    aaa
        radius-server-policy "rad-serv-pol-1" create
            servers
```

```
                router "Base"
                source-address 192.0.2.1
                server 1 name "radius-172.16.1.2"
            exit
        exit
    exit
exit

configure
    subscriber-mgmt
        authentication-policy "radius-1" create
            description "Radius authentication policy"
            password vsecret1
            radius-authentication-server
                source-address 192.0.2.1
            exit
            radius-server-policy "rad-serv-pol-1"
        exit
        radius-accounting-policy "nat-accounting" create
            update-interval 5
            include-radius-attribute
                mac-address
                nat-port-range
                subscriber-id
            exit
            radius-accounting-server
                source-address 192.0.2.1
                router "Base"
                server 1 address 172.16.1.2 secret vsecret1
            exit
        exit
    exit
exit
```

The initial service and ESM configuration is given below for VPRN 2. This VPRN contains subscriber interface sub-int-2 with group interface group-int-pppoe-cpe, in routed CO. There is one SAP under the group interface, 1/1/1:111. The subscribers are PPP over Ethernet, and are authenticated through RADIUS.

```
configure
    service
        vprn 2 customer 1 create
            route-distinguisher 64496:2
            subscriber-interface "sub-int-2" create
                address 10.0.1.254/24
                group-interface "group-int-pppoe-cpe" create
                    authentication-policy "radius-1"
                    sap 1/1/1:111 create
                        sub-sla-mgmt
                            def-sub-profile "sub-profile-nat"
                            def-sla-profile "sla-profile-nat"
                            sub-ident-policy "sub-ident-nat"
                            multi-sub-sap 10
                            no shutdown
                        exit
                    exit
                    pppoe
```

```
                        policy "pppoe-nat"
                        session-limit 10
                        sap-session-limit 10
                        no shutdown
                    exit
                exit
            exit
            no shutdown
        exit
    exit
exit
```

Successful Creation of a NAT Subscriber

When a subscriber-host is created and a NAT-policy is defined, its inside IP address should fall within the L2-aware range (see next section). If this is not the case, the subscriber-host creation will fail. As mentioned before, NAT is not supported on static-hosts. The subscriber-host creation is shown in Figure 68. If the subscriber-host does not get an inside or outside IP address, it would not be able to communicate with any servers on the outside.

Figure 68 **Subscriber-Host Creation Flow**

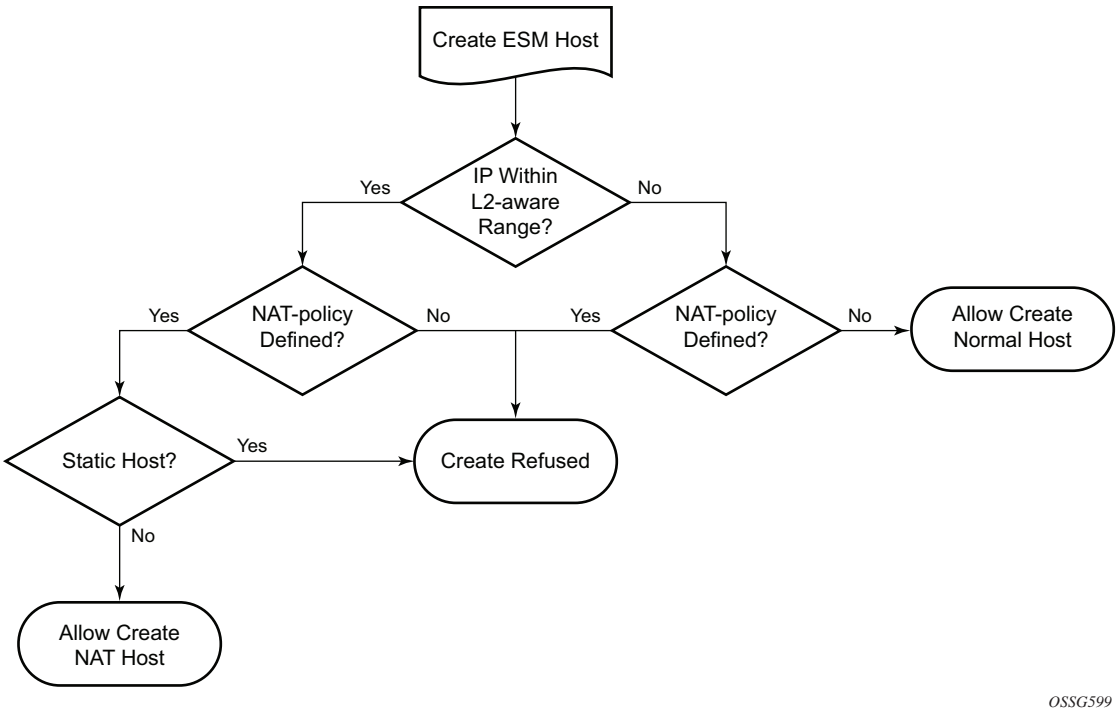

NAT Inside Configuration

The residential gateway subnets which are to be NATed need to be configured in the nat inside L2-aware context. In this chapter, all subscribers belonging to VPRN 1 will be allocated an address in subnet 10.0.0.0/24. All subscribers belonging to VPRN 2 will be allocated an IP address in subnet 10.0.1.0/24. Hosts in these services and within these subnets will be subject to L2-aware NAT if they have the correct nat-policy in their subscriber profile.

The address configured here is the local IP address of the system, and is typically used as next hop for a L2-aware host. Hosts connected to the inside service can ARP for this address. To verify connectivity, a host can ping this address.

```
configure
    service
        vprn 1 customer 1 create
            nat
                inside
                    l2-aware
                        address 10.0.0.254/24
                    exit
                exit
            exit
        exit
    exit
exit

configure
    service
        vprn 2 customer 1 create
            nat
                inside
                    l2-aware
                        address 10.0.1.254/24
                    exit
                exit
            exit
        exit
    exit
exit
```

NAT Outside Configuration

The NAT outside pool needs to be configured on the VPRN facing the outside world, which is the public Internet. The NAT outside pool controls the NAT type, which in this case is L2-aware, and the NAT group to send the traffic to.

These addresses will be used to as source addresses for all packets in the upstream direction (toward the public Internet) and as destination address for all packets in the downstream direction.

```
configure
    service
        vprn 1 customer 1 create
            nat
                outside
                    pool "nat-outside-pool-1" nat-group 1 type l2-aware create
                        port-reservation blocks 128
                        address-range 10.255.0.1 10.255.0.10 create
                        exit
                        no shutdown
                    exit
                exit
            exit
         exit
    exit
exit
```

The port-reservation command specifies the number of port-blocks (blocks of consecutive usable port numbers) per IP address. In this configuration, each public IP address is subdivided into 128 port blocks which can be used for NAT, that results in 504 public ports per block.

Binding Inside NAT, Outside NAT and ESM Host

In order to bind the inside part of the NAT with the outside part a nat-policy needs to be created, in the service nat context. The outside nat-pool, which is associated with the VPRN instance in this example, and outside IP addresses, are configured under the nat-policy.

```
configure
    service
        nat
            nat-policy "nat-l2aware-vprn1" create
                pool "nat-outside-pool-1" router 1
                timeouts
                    icmp-query min 4
                exit
            exit
        exit
    exit
exit
```

This nat-policy is then associated to the different subscribers by means of the subscriber profile. The ESM host originated traffic will be diverted to the NAT device.

```
configure
    subscriber-mgmt
        sub-profile "sub-profile-nat" create
            nat-policy "nat-l2aware-vprn1"
        exit
    exit
exit
```

The nat-policy also controls the following parameters:

Filtering — Two filtering modes are available, endpoint-independent (default configuration) and address-and-port dependent filtering. The filtering behavior will control which upstream packets are transmitted (based on the existing sessions). If endpoint-independent filtering is configured, any outside host/port can use mappings the NAT has created to send traffic to the inside. If address-and-port-dependent filtering is selected, only packets from the same destination and port which created the mapping will be processed.

Port limits — A number of ports can be reserved for prioritized sessions. A session is considered as a priority-session depending on its forwarding class. High and low watermarks can be configured to trigger alarms based on the port usage.

The reserved resources mean that if the resources get down to the level that there is only the reserved amount left, this leftover can only be used by priority sessions, not taking into account the amount of priority sessions already set up at that point.

Example: By default each host is assigned 504 outside ports. 100 of these ports can be reserved for the EF and H1 forwarding classes. As soon as any given host reaches 404 utilized outside ports, the remaining 100 will only be used for EF or H1 sessions.

Priority sessions — The forwarding classes for which the sessions should be prioritized in terms of port or session assignment can be configured here. Multiple forwarding classes can be configured.

Session limits — A maximum number of sessions can be configured for each subscriber associated with this nat-policy. A number of sessions can be reserved for prioritized sessions. Sessions are prioritized based on forwarding class. High and low watermarks can be configured to trigger alarms based on the session usage.

Note:

- The reserved sessions and reserved ports are not the same. A user can have many applications contacting the same destination. Many different source ports will be used, therefore many different outside ports. A user can have one application contacting many different destinations. The same source port, but many different destination IP addresses will be used. Only one outside port is consumed, but many sessions exist.
- It is possible to configure a reserved ports session-limit on the nat-group as well. In case both per Layer 2 aware host and per nat-group limits are configured the most restrictive will be enforced.

Timeouts — Several timeouts can be configured.

- icmp-query: Timeout applied to an ICMP query session.
- tcp-established: The idle timeout applied to a TCP session in the established state.
- tcp-syn: The timeout applied to a TCP session in the SYN state.
- tcp-time-wait: Time-wait assassination is enabled by default to quickly remove TCP mappings in the time-wait state.
- tcp-transitory: The idle timeout applied to a TCP session in a transitory state. TCP transition between SYN and Open.
- udp: All udp streams (with exceptions of udp-initial and udp-dns).
- udp-initial: UDP mapping timeout applied to new sessions. Applicable when only 1 UDP packet is sent.
- udp-dns: Only traffic to destination UDP port 53

Advanced Topics

RADIUS Accounting

RADIUS accounting is extended with a new attribute, nat-port-range, reporting the NAT port range in use by the subscriber. In order to configure RADIUS accounting, first the RADIUS accounting policy must be created.

```
configure
    subscriber-mgmt
        radius-accounting-policy "nat-accounting" create
            update-interval 5
            include-radius-attribute
                mac-address
```

```
                nat-port-range
                subscriber-id
            exit
            radius-accounting-server
                source-address 192.0.2.1
                router "Base"
                server 1 address 172.16.1.2 secret vsecret1
            exit
        exit
    exit
exit
```

The configuration specifies which attributes to include in the radius accounting messages toward the configured server. The update interval is specified in minutes. Every 5 minutes an update will be sent to the RADIUS accounting server.

Then this RADIUS accounting policy must be attached to the subscriber profile.

```
configure
    subscriber-mgmt
        sub-profile "sub-profile-nat"
            nat-policy "nat-l2aware-vprn1"
            radius-accounting
                policy "nat-accounting"
            exit
        exit
    exit
exit
```

Hardware Resource Monitoring

It is possible to define watermarks to monitor the actual usage of sessions and/or ports. For each watermark, a high and a low value have to be set (as a percentage). Once the high value is reached, a notification will be sent. As soon as the usage drops below the low watermark, another notification (trap) will be sent.

Watermarks can be defined on nat-group, pool and policy level.

NAT-group

Watermarks can be configured to monitor the total number of sessions on an MDA.

```
configure
    isa
        nat-group 1
            session-limits
                watermarks high 90 low 80
            exit
        exit
```

```
    exit
exit
```

Pool

Watermarks can be configured to monitor the total number of blocks in use in a pool.

```
configure
    service
        vprn 1 customer 1
            nat
                outside
                    pool "nat-outside-pool-1"
                        watermarks high 90 low 80
                    exit
                exit
            exit
         exit
    exit
exit
```

Policy

In the policy it is possible to define watermarks on session and port usage. The usage per subscriber will be monitored.

```
configure
    service
        nat
            nat-policy "nat-l2aware-vprn1"
                port-limits
                    watermarks high 90 low 80
                exit
                session-limits
                    watermarks high 90 low 80
                exit
            exit
        exit
    exit
exit
```

Outside IP Address Range Management

From an operational point of view, it may be required to unprovision an outside IP address range. To that end, the **drain** has been introduced. If configured, no new sessions will be set up using this address-range. Existing mappings will cease to exist only when the session ends (tcp fin, fin ack, ack) or other timeout mechanism.

```
configure
    service
        vprn 1
            nat
                outside
                    pool "nat-outside-pool-1" nat-group 1
                        address-range 10.255.0.1 10.255.0.10
                            drain
                        exit
                        no shutdown
                    exit
                exit
            exit
         exit
    exit
exit
```

When all sessions have drained the address-range can be unprovisioned.

Quality of Service

NAT is fully transparent in terms of quality of service. The quality of service is determined on ingress into the service router. A forwarding class is assigned to each packet and is retained throughout the whole router.

For L2-aware NAT, a virtual port exists on the carrier IOM, nat-in-l2. This port is modeled as a network port with per FC queues both on ingress and egress. On network-ingress per destination queues are implemented, making sure head of line blocking cannot happen.

Operation

The MS-ISA card should be in operational up state.

```
*A:P1# show mda

===============================================================================
MDA Summary
===============================================================================
Slot  Mda   Provisioned Type                          Admin     Operational
                Equipped Type (if different)         State     State
-------------------------------------------------------------------------------
1     1     m4-10gb-xp-xfp                           up        up
      2     isa-bb                                   up        up
                isa-ms
===============================================================================
*A:P1#
```

The NAT group should be configured, with at least one pool of outside IP addresses associated with it. In this example, redundancy is active-standby, but active-active is also possible.

```
*A:P1# show isa nat-group 1

===============================================================================
ISA NAT Group 1
===============================================================================
Description                 : L2 aware NAT group
Admin state                 : inService
Operational state           : inService
Degraded                    : false
Redundancy                  : active-standby
Active MDA limit            : 1
Failed MDA limit            : 0

-------------------------------------------------------------------------------
NAT specific information for ISA group 1
-------------------------------------------------------------------------------
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            : 04/28/2016 12:20:16
-------------------------------------------------------------------------------
===============================================================================

===============================================================================
ISA Group 1 members
===============================================================================
Group Member     State          Mda  Addresses  Blocks     Se-% Hi Se-Prio
-------------------------------------------------------------------------------
1     1          active         1/2  10         1280       < 1  N  0
-------------------------------------------------------------------------------
No. of members: 1
===============================================================================
*A:P1#
```

The following table describes the **show isa nat-group** output fields.

Table 8 **Show isa nat-group Output Field Descriptions**

Field	Description
Group	The group-id
Member	All members will be listed with associated parameters
State	The operational state of each member
MDA	The MDA position of the member

Table 8 **Show isa nat-group Output Field Descriptions (Continued)**

Field	Description (Continued)
Addresses	The number of outside IP addresses assigned to the member
Blocks	The number of allocated port-blocks
Se-%	The actual session usage in percentage
Hi	High watermark reached (Y/N)
Se-Prio	The configured number of priority sessions

```
*A:P1# show isa nat-group 1 associations

===============================================================================
ISA NAT Group 1 pool associations
===============================================================================
Pool                                    Router
-------------------------------------------------------------------------------
nat-outside-pool-1                      vprn1
-------------------------------------------------------------------------------
No. of pools: 1
===============================================================================
No associated router instances found.
```

The subscriber-hosts should be created correctly.

```
*A:P1# show service id 1 subscriber-hosts

============================================================
Subscriber Host table
============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID Origin       Fwding State
------------------------------------------------------------
1/1/1:81               ipoe-sub-00:0c:29:9d:10:2d
  10.0.0.2
    00:0c:29:9d:10:2d    N/A       DHCP         Fwding
1/1/1:82               ipoe-sub-00:0c:29:34:cc:74
  10.0.0.2
    00:0c:29:34:cc:74    N/A       DHCP         Fwding
------------------------------------------------------------
Number of subscriber hosts : 2
============================================================
*A:P1#

*A:P1# show service id 2 subscriber-hosts

============================================================
Subscriber Host table
============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID Origin       Fwding State
```

```
-------------------------------------------------------------
1/1/1:111              pppoe-sub-00:0c:29:1d:44:34
  10.0.1.2
    00:0c:29:1d:44:34    1         IPCP         Fwding
-------------------------------------------------------------
Number of subscriber hosts : 1
=============================================================
*A:P1#
```

The associated L2-aware NAT subscriber-hosts are visible from CLI. The associated group, member and ports can be viewed using this command.

```
*A:P1# show service nat l2-aware-subscribers

===============================================================================
Layer-2-Aware NAT subscribers
===============================================================================

Subscriber                          : ipoe-sub-00:0c:29:34:cc:74
-------------------------------------------------------------------------------
ISA NAT group                       : 1
ISA NAT group member                : 1
UPnP policy                         : (None)
Default NAT policy                  : nat-l2aware-vprn1

Policy                              : nat-l2aware-vprn1
Outside router                      : vprn1
Outside IP                          : 10.255.0.1
Ports                               : 1056-1087

Subscriber                          : ipoe-sub-00:0c:29:9d:10:2d
-------------------------------------------------------------------------------
ISA NAT group                       : 1
ISA NAT group member                : 1
UPnP policy                         : (None)
Default NAT policy                  : nat-l2aware-vprn1

Policy                              : nat-l2aware-vprn1
Outside router                      : vprn1
Outside IP                          : 10.255.0.10
Ports                               : 1024-1055

Subscriber                          : pppoe-sub-00:0c:29:1d:44:34
-------------------------------------------------------------------------------
ISA NAT group                       : 1
ISA NAT group member                : 1
UPnP policy                         : (None)
Default NAT policy                  : nat-l2aware-vprn1

Policy                              : nat-l2aware-vprn1
Outside router                      : vprn1
Outside IP                          : 10.255.0.3
Ports                               : 1056-1087

-------------------------------------------------------------------------------
No. of subscribers: 3
===============================================================================
*A:P1#
```

The active subscribers can be shown using following command.

```
*A:P1# show service active-subscribers
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber ipoe-sub-00:0c:29:34:cc:74 (sub-profile-nat)
-------------------------------------------------------------------------------
NAT Policy: nat-l2aware-vprn1
Outside IP: 10.255.0.1 (vprn1)
Ports     : 1056-1087

-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:1/1/1:82 - sla:sla-profile-nat
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
10.0.0.2
                00:0c:29:34:cc:74    N/A            DHCP         1          Y
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Subscriber ipoe-sub-00:0c:29:9d:10:2d (sub-profile-nat)
-------------------------------------------------------------------------------
NAT Policy: nat-l2aware-vprn1
Outside IP: 10.255.0.10 (vprn1)
Ports     : 1024-1055

-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:1/1/1:81 - sla:sla-profile-nat
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
10.0.0.2
                00:0c:29:9d:10:2d    N/A            DHCP         1          Y
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Subscriber pppoe-sub-00:0c:29:1d:44:34 (sub-profile-nat)
-------------------------------------------------------------------------------
NAT Policy: nat-l2aware-vprn1
Outside IP: 10.255.0.3 (vprn1)
Ports     : 1056-1087

-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:1/1/1:111 - sla:sla-profile-nat
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
10.0.1.2
                00:0c:29:1d:44:34    PPP 1          IPCP         2          Y
-------------------------------------------------------------------------------
```

```
-------------------------------------------------------------------------------
Number of active subscribers : 3
-------------------------------------------------------------------------------
*A:P1#
```

Traffic arriving on the node from the outside world should be routed to the correct MS-ISA card (=NAT device).

The route table of VPRN 1 indicates that all traffic destinedto the publicly visible IP addresses are routed to mda 1/2. In other words, all packets coming from the outside toward the 10.255.0.1 to 10.255.0.10 IP addresses are sent to the NAT device.

```
*A:P1# show router 1 route-table

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.0.0.0/24                                     Local   Local     22h46m55s  0
       sub-int-1                                                       0
10.0.1.0/24                                     Local   Local     22h46m55s  0
       sub-int-1                                                       0
10.11.11.1/32                                   Local   Local     22h46m55s  0
       int-TEST                                                        0
10.255.0.1/32                                   Remote  NAT       22h34m44s  0
       NAT outside to mda 1/2                                          0
10.255.0.2/31                                   Remote  NAT       22h34m44s  0
       NAT outside to mda 1/2                                          0
10.255.0.4/30                                   Remote  NAT       22h34m44s  0
       NAT outside to mda 1/2                                          0
10.255.0.8/31                                   Remote  NAT       22h34m44s  0
       NAT outside to mda 1/2                                          0
10.255.0.10/32                                  Remote  NAT       22h34m44s  0
       NAT outside to mda 1/2                                          0
172.16.2.0/30                                   Local   Local     22h46m55s  0
       int-PE-1-servers                                                0
-------------------------------------------------------------------------------
No. of Routes: 9
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
*A:P1#

*A:P1# show router 2 route-table

===============================================================================
Route Table (Service: 2)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.0.1.0/24                                     Local   Local     22h43m06s  0
```

```
       sub-int-2                                                              0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

Individual sessions are viewed through a tools dump command.

```
*A:P1# tools dump nat sessions

===============================================================================
Matched 3 sessions on Slot #1 MDA #2
===============================================================================
Owner               : L2-aware Subscr ipoe-sub-00:0c:29:9d:10:2d
Policy              : nat-l2aware-vprn1
FlowType            : ICMP              Timeout (sec)       : 194
Inside IP Addr      : 10.0.0.2
Inside Identifier   : 0
Outside IP Addr     : 10.255.0.10
Outside Identifier  : 1048
Foreign IP Addr     : 172.16.2.2
Nat Group           : 1
Nat Group Member    : 1
-------------------------------------------------------------------------------
Owner               : L2-aware Subscr pppoe-sub-00:0c:29:1d:44:34
Policy              : nat-l2aware-vprn1
FlowType            : ICMP              Timeout (sec)       : 192
Inside IP Addr      : 10.0.1.2
Inside Identifier   : 0
Outside IP Addr     : 10.255.0.3
Outside Identifier  : 1086
Foreign IP Addr     : 172.16.2.2
Nat Group           : 1
Nat Group Member    : 1
-------------------------------------------------------------------------------
Owner               : L2-aware Subscr ipoe-sub-00:0c:29:34:cc:74
Policy              : nat-l2aware-vprn1
FlowType            : ICMP              Timeout (sec)       : 204
Inside IP Addr      : 10.0.0.2
Inside Identifier   : 0
Outside IP Addr     : 10.255.0.1
Outside Identifier  : 1076
Foreign IP Addr     : 172.16.2.2
Nat Group           : 1
Nat Group Member    : 1
-------------------------------------------------------------------------------
===============================================================================
*A:P1#
```

The resources on the MS-ISA can also be viewed through a **tools dump** command.

```
*A:P1# tools dump nat isa resources mda 1/2

Resource Usage for Slot #1 Mda #2:
```

```
                            | Total        | Allocated    | Free          | Limit
----------------------------+--------------+--------------+---------------+------
                      Flows |       196608 |            3 |        196605 |   N/A
                   Policies |         1024 |            1 |          1023 |   N/A
     Port-ranges configured |       524288 |            3 |        524285 |  100%
           Port-ranges used |            3 |            3 |             0 |  100%
       Port-ranges retained |            0 |            0 |             0 |  100%
                      Ports |   1006632960 |            3 |    1006632957 |  100%
                IP-addresses |        65536 |            3 |         65533 |  100%
          Large-scale hosts |         8192 |            0 |          8192 |  100%
   Subscriber-cache entries |         8192 |            0 |          8192 |   N/A
       L2-aware subscribers |         2048 |            3 |          2045 |  100%
             L2-aware hosts |         2048 |            3 |          2045 |  100%
             Delayed ICMP's |          200 |            0 |           200 |   N/A
                ALG session |        24576 |            0 |         24576 |   N/A
    Upstream fragment lists |         2048 |            0 |          2048 |   N/A
  Downstream fragment lists |         1024 |            0 |          1024 |   N/A
    Upstream fragment holes |         8192 |            0 |          8192 |   N/A
  Downstream fragment holes |         4096 |            0 |          4096 |   N/A
     Upstream fragment bufs |         2048 |            0 |          2048 |   N/A
   Downstream fragment bufs |         1024 |            0 |          1024 |   N/A
        Dormant subscribers |            0 |            0 |             0 |   N/A
              UPnP mappings |         1024 |            0 |          1024 |   N/A
              UPnP sessions |          100 |            0 |           100 |   N/A
    One-to-one IP-addresses |         8192 |            0 |          8192 |  100%
 Flowlog destinations set 0 |            2 |            0 |             2 |   N/A
 Flowlog destinations set 1 |            2 |            0 |             2 |   N/A
 Flowlog destinations set 2 |            1 |            0 |             1 |   N/A
      Flowlog packets set 0 |          128 |            0 |           128 |   N/A
      Flowlog packets set 1 |          128 |            0 |           128 |   N/A
      Flowlog packets set 2 |          128 |            0 |           128 |   N/A

*A:P1#
```

The alarm configuration can be verified for the NAT related traps.

```
*A:P1# show log event-control "nat"
===============================================================================
Log Events
===============================================================================
Application
 ID#    Event Name                          P   g/s     Logged     Dropped
-------------------------------------------------------------------------------
   2001 tmnxNatPlL2AwBlockUsageHigh         WA  thr          0           0
   2002 tmnxNatIsaMemberSessionUsageHigh    WA  thr          0           0
   2003 tmnxNatPlLsnMemberBlockUsageHigh    WA  thr          0           0
   2007 tmnxNatL2AwSubIcmpPortUsageHigh     WA  thr          0           0
   2008 tmnxNatL2AwSubUdpPortUsageHigh      WA  thr          0           0
   2009 tmnxNatL2AwSubTcpPortUsageHigh      WA  thr          0           0
   2010 tmnxNatL2AwSubSessionUsageHigh      WA  thr          0           0
   2012 tmnxNatPlBlockAllocationLsn         MI  sup          0           0
   2013 tmnxNatPlBlockAllocationL2Aw        MI  sup          0          23
   2014 tmnxNatResourceProblemDetected      MI  thr          0           0
   2015 tmnxNatResourceProblemCause         MI  thr          0           0
   2016 tmnxNatPlAddrFree                   MI  sup          0           0
   2017 tmnxNatPlLsnRedActiveChanged        WA  thr          0           0
   2018 tmnxNatPcpSrvStateChanged           MI  thr          0           0
```

```
   2020 tmnxNatMdaActive                  MI  thr          1           0
   2021 tmnxNatLsnSubBlksFree             MI  sup          0           0
   2022 tmnxNatDetPlcyChanged             MI  thr          0           0
   2023 tmnxNatMdaDetectsLoadSharingErr   MI  thr          0           0
   2024 tmnxNatIsaGrpOperStateChanged     MI  thr          2           0
   2025 tmnxNatIsaGrpIsDegraded           MI  thr          1           0
   2026 tmnxNatLsnSubIcmpPortUsgHigh      WA  thr          0           0
   2027 tmnxNatLsnSubUdpPortUsgHigh       WA  thr          0           0
   2028 tmnxNatLsnSubTcpPortUsgHigh       WA  thr          0           0
   2029 tmnxNatLsnSubSessionUsgHigh       WA  thr          0           0
   2030 tmnxNatInAddrPrefixBlksFree       MI  sup          0           0
   2031 tmnxNatFwd2EntryAdded             MI  sup          0           0
   2032 tmnxNatDetPlcyOperStateChanged    MI  thr          0           0
   2033 tmnxNatDetMapOperStateChanged     MI  thr          0           0
   2034 tmnxNatFwd2OperStateChanged       WA  thr          0           0
   2035 tmnxNatVrtrOutDnatOnlyRoutesHigh  WA  thr          0           0
===============================================================================
*A:P1#
```

RADIUS accounting information can be verified using the **debug router radius detail** command.

```
1 2016/04/29 09:11:45.46 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Accounting-Request(4) 172.16.1.2:1813 id 248 len 316 vrid 1 pol nat-accounting
    STATUS TYPE [40] 4 Interim-Update(3)
    NAS IP ADDRESS [4] 4 192.0.2.1
    SESSION ID [44] 71 ipoe-sub-00:0c:29:9d:10:2d@1/1/1:81@sla-profile-nat_2016/
04/28 12:21:07
    SESSION TIME [46] 4 75038
    EVENT TIMESTAMP [55] 4 1461921105
    VSA [26] 193 Alcatel(6527)
      SUBSC ID STR [11] 26 ipoe-sub-00:0c:29:9d:10:2d
      SUBSC NAT PORT RANGE [121] 48 10.255.0.10 1024-1055 router 1 nat-l2aware-v
prn1
      CHADDR [27] 17 00:0c:29:9d:10:2d
      INPUT_INPROF_OCTETS_64 [19] 10 0x00010000000000000000
      INPUT_OUTPROF_OCTETS_64 [20] 10 0x0001000000000029b40
      INPUT_INPROF_PACKETS_64 [23] 10 0x00010000000000000000
      INPUT_OUTPROF_PACKETS_64 [24] 10 0x000100000000000009d0
      OUTPUT_INPROF_OCTETS_64 [21] 10 0x00010000000000000444
      OUTPUT_OUTPROF_OCTETS_64 [22] 10 0x00010000000000028154
      OUTPUT_INPROF_PACKETS_64 [25] 10 0x0001000000000000000e
      OUTPUT_OUTPROF_PACKETS_64 [26] 10 0x000100000000000009b7

"

2 2016/04/29 09:11:45.46 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Accounting-Response(5) id 248 len 20 from 172.16.1.2:1813 vrid 1 pol nat-accou
nting

"
```

Conclusion

L2-aware NAT provides IPv4 NAT services to ESM subscribers.

This chapter shows the configuration of L2-aware NAT together with the associated show outputs which can be used verify and troubleshoot it.

NAT Stateless Dual-Homing

This chapter describes NAT stateless dual-homing.

Topics in this chapter include:

Applicability

The information and configuration in this chapter is based on SR OS release 14.0.R4.

This feature is applicable to:

- 7750 SR-7/12 equipped with:
  - MS-ISM
  - MS-ISM-E
  - MS-ISA2 combo cards
  - MS-ISA2-E combo cards
  - IOM2 or IOM3-XP/-B/-C housing:
    - MS-ISA
    - MS-ISA-E
  - IOM4-e/IOM4-e-B housing:
    - MS-ISA2
    - MS-ISA2-E
- 7750 SR-1e/2e/3e equipped with IOM-e housing:
  - MS-ISA2
  - MS-ISA2-E
- 7450 ESS-7/12 in mixed mode equipped with:
  - MS-ISM
  - MS-ISA2 combo cards
  - IOM2 or IOM3-XP/-B/-C housing:

- MS-ISA
- MS-ISA-E

– IOM4-e/IOM4-e-B housing:

- MS-ISA2
- MS-ISA2-E

Overview

With the IPv4 address space almost consumed, many operators are deploying Network Address Translation (NAT) at centralized or semi-centralized points in their IP/MPLS networks. The NAT function is implemented using Carrier Grade NAT (CGN) nodes, which typically support tens of thousands of clients/subscribers. Therefore, a failure of one of these nodes would be considered a significant event.

Many operators consider a stateful failover mechanism between CGN nodes to be too demanding with regard to control plane requirements and state synchronization of NAT bindings. A reasonable compromise appears to be a stateless failover mechanism, capable of providing failover between geo-redundant CGN devices, but with manageable control plane implications.

This chapter describes the NAT stateless dual-homing feature supported in SR OS. NAT stateless dual-homing is supported for Large-Scale NAT and NAT64, and this chapter describes how both can be supported, either independently or together.

Example Topology

The topology shown in Figure 69 is an example of the use of NAT stateless dual-homing. PE-2, PE-4, RR-10, CGN-1, and CGN-2 form part of Autonomous System 64496 and run IS-IS level 2 and LDP. PE-2, PE-4, CGN-1, and CGN-2 are clients of Route Reflector RR-10 and peer with the VPN-IPv4 and VPN-IPv6 address families.

IAR-1 acts as an Internet Service Provider (ISP) edge router belonging to AS 65535, and provides Internet access to AS 65596. CGN-1 and CGN-2 are configured with a VPRN (1000) that serves as a NAT outside routing VPN Routing and Forwarding (VRF) instance. In this VRF, both CGN-1 and CGN-2 peer with IAR-1 in EBGP for the IPv4 address family. CGN-1 and CGN-2 advertise the relevant NAT outside pools to IAR-1, and IAR-1 advertises a default route to both CGN-1 and CGN-2. IAR-1 has IP connectivity to a web server at 10.1.128.43/24.

CE-2 is connected to PE-2 and is part of VPRN 200. CE-2 has an IPv4 host (172.31.102.3) that serves to test NAT44 connectivity to the web server, and an IPv6 host (2001:DB8:4001:200::3) that serves to test NAT64 connectivity to the web server.

Figure 69 **Example Topology**

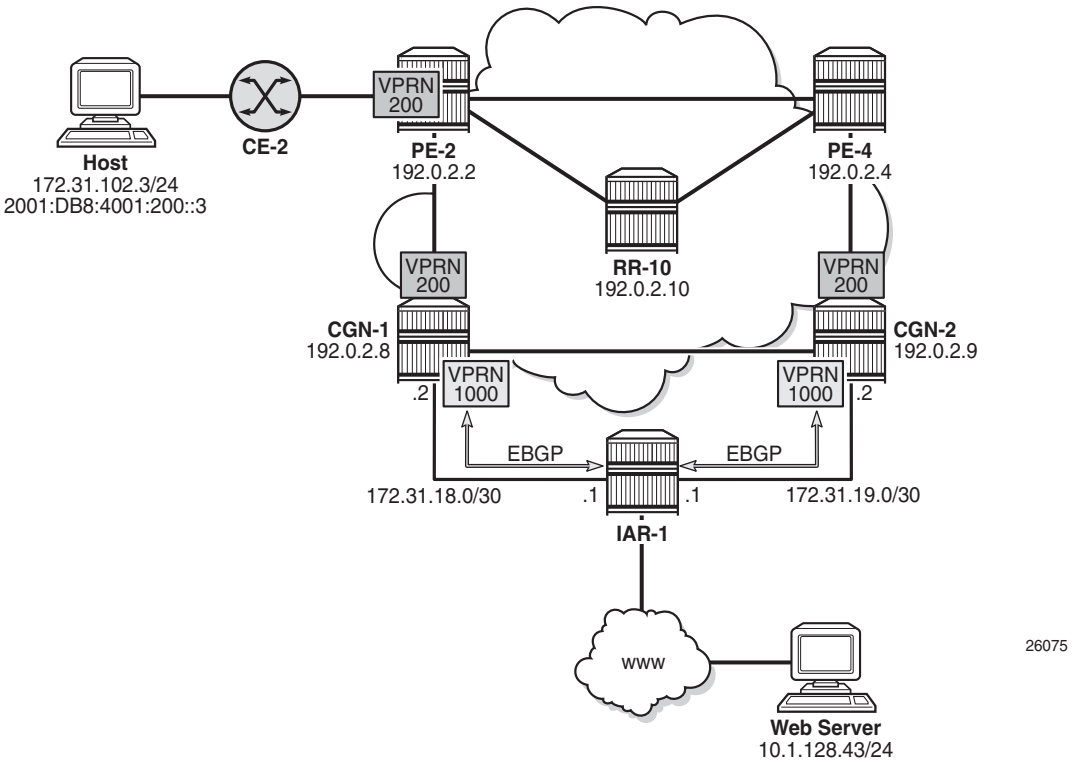

Configuration

To support NAT functionality, some form of Integrated Services Adapter (ISA) or Integrated Services Module (ISM) is required, as listed in the Applicability section. In this example, CGN-1 and CGN-2 both have a single MS-ISA card, which is configured as MDA type **isa-bb** and is housed in a carrier IOM of type IOM3-XP-B, as follows:

```
configure
    card 1
        card-type iom3-xp-b
        mda 2
            mda-type isa-bb
            no shutdown
        exit
```

```
        no shutdown
    exit
```

The MS-ISA is then configured to become a member of one or more **nat-groups**. Up to fourteen MS-ISAs can be configured to belong to up to four NAT groups. When more than one MS-ISA is configured in a NAT group, the MS-ISAs can work in active-standby mode or active-active mode.

In active-standby mode, one or more MS-ISAs act as standby, and in normal operation, are idle. If an active MS-ISA fails, one standby MS-ISA accepts the traffic from the failed card. In the active-standby scenario, the mapping between failed card and standby card can always be considered to be 1:1. In active-active mode, all of the MS-ISAs in the NAT group are active, and if one MS-ISA in the group fails, the load is distributed across the remaining active MS-ISAs in the group.

The default setting is active-standby. In both active-standby and active-active modes, the dynamically created NAT bindings are not synchronized between cards. Therefore, a failover will cause an interruption in traffic until the NAT bindings are re-initiated by the clients/subscribers behind the NAT.

The NAT group also requires that an **active-mda-limit** is configured, as follows, which allows the operator to specify how many MDAs (MS-ISAs) will be active in the group. Any operational MDAs above this configured limit will be considered spare MDAs. Finally, the **nat-group** must be placed in a **no shutdown** state.

```
configure
    isa
        nat-group 1 create
            active-mda-limit 1
            mda 1/2
            no shutdown
        exit
    exit
```

NAT Outside Context

The NAT *outside* function is responsible for creation of the NAT bindings, using outside IP addresses defined in outside pools (together with their associated ports), and for advertising the address ranges in those pools to upstream routers. The NAT stateless dual-homing redundancy mechanism is based on ownership of an outside pool, where each member of a redundant pair can assume either an active (master) or standby role for an outside pool. This active/standby role is determined by the presence of a *monitor* prefix. Both CGN nodes of a redundant pair implement the following:

1. Advertise a unique route into the routing instance that the NAT outside function resides in. This is known as the export route and may be advertised into the Global Routing Table (GRT) or a VPRN instance. For example, CGN-1 advertises (exports) prefix P1 while CGN-2 advertises prefix P2.
2. Check for the presence of a configured route in the routing instance that the NAT outside function resides in. This is known as the monitor route. Continuing the preceding example, CGN-1 monitors prefix P2 while CGN-2 monitors prefix P1.

Therefore, the export route of CGN-1 becomes the monitor route of CGN-2, and the export route of CGN-2 becomes the monitor route of CGN-1. The redundancy mechanism thereafter checks the (virtual) routing table of the NAT outside function for the presence of the monitor route. If it is not present, the redundancy state for the pool is set to active and the following occurs (subject to routing policy):

1. The redundancy export route is populated in the NAT outside routing instance and advertised externally in the relevant routing instance.
2. The outside pool address is populated in the NAT outside routing instance and advertised externally in the relevant routing instance.
3. Routes that need to become active in any associated NAT inside routing instances, to attract traffic to the active CGN node, are populated in the relevant routing tables. For example, NAT64 translator routes and/or NAT44 steering routes, both of which are described later in this chapter, are populated.

If the monitor route is present, the redundancy state for the pool is set to standby and the following occurs:

1. The redundancy export route is not populated in the NAT outside routing instance and is, therefore, not eligible to be advertised externally.
2. The outside pool address is not populated in the NAT outside routing instance and is, therefore, not eligible to be advertised externally.
3. Routes that need to become active in any associated NAT inside routing instances, to attract traffic to the active CGN node, are not populated in the relevant routing tables.

There are no configurable options for selection of active/standby CGN nodes. The status of a node is based on the presence of the monitor route. In the event of a collision during redundancy startup, hardcoded debounce timers ensure that only a single CGN node is selected as active.

The first of two following examples shows the configuration of VPRN 1000 (the NAT outside VRF) at CGN-2, with the second configuration showing the **vrf-export** and **vrf-import** policy statements. For advertising and importing routes from/to the VRF, there are two requirements: to advertise the export redundancy route and to import the monitor route. This is the purpose of the **"vrf1000-export"** and **"vrf1000-import"** policy statements. Other VPRN parameters are generic and, therefore, not discussed here.

The VPRN contains an interface, "to-IAR-1", which has an associated EBGP peering session to IAR-1. The export policy under the BGP **neighbor** context will be described later in this chapter. That policy contains sufficient logic to advertise the NAT outside pools.

The NAT pools are configured in the **nat outside** context. The first configuration provides an example of a single pool, "4-to-4", which will be used for NAT44. As well as a name, the pool requires association with a NAT group, and definition of the **type** of NAT that will be configured; in this case, **large-scale**. The **mode** of the pool is set to **napt** to indicate N:1 NAT, and the **address-range** that will be used for outside addressing is 10.1.4.1 to 10.1.4.254.

The relevant and required parameters for stateless dual-homing are configured in the **redundancy** context. In this example, CGN-2 exports prefix 192.168.0.249/32 and monitors prefix 192.168.0.248/32. Conversely, CGN-1 exports prefix 192.168.0.248/32 and monitors prefix 192.168.0.249/32. The redundancy node and the pool must be placed into a **no shutdown** state.

```
*A:CGN-2#
configure
    service
        vprn 1000 customer 1 create
            vrf-import "vrf1000-import"
            vrf-export "vrf1000-export"
            autonomous-system 64496
            route-distinguisher 64496:1000
            auto-bind-tunnel
                resolution any
            exit
            interface "to-IAR-1" create
                address 172.31.19.2/30
                sap 1/1/2:200 create
                exit
            exit
            aggregate 10.1.4.0/24 summary-only
            bgp
                group "EBGP"
                    family ipv4
                    peer-as 65535
                    split-horizon
                    neighbor 172.31.19.1
                        authentication-key <password>
                        export "vrf1000-ebgp-export"
                    exit
```

```
                    exit
                    no shutdown
                exit
                nat
                    outside
                        pool "4-to-4" nat-group 1 type large-scale create
                            redundancy
                                export 192.168.0.249/32
                                monitor 192.168.0.248/32
                                no shutdown
                            exit
                            mode napt
                            address-range 10.1.4.1 10.1.4.254 create
                            exit
                            no shutdown
                        exit
                    exit
                exit
                service-name "NAT-Outside"
                no shutdown
            exit
        exit

    *A:CGN-2#
    configure
       router
            policy-options
                begin
                prefix-list "vrf1000-nat-export"
                    prefix 192.168.0.249/32 exact
                exit
                community "vrf1000-export" members "target:64496:1000"
                community "vrf1000-import" members "target:64496:1000"
                policy-statement "vrf1000-export"
                    entry 10
                        from
                            protocol nat
                            prefix-list "vrf1000-nat-export"
                        exit
                        to
                            protocol bgp-vpn
                        exit
                        action accept
                            community add "vrf1000-export"
                        exit
                    exit
                exit
                policy-statement "vrf1000-import"
                    entry 10
                        from
                            community "vrf1000-import"
                        exit
                        action accept
                        exit
                    exit
                    default-action drop
                    exit
                exit
```

After the redundancy node and pool are placed into a **no shutdown** state, the master and standby can be elected, based on the previously described criteria. In this example, CGN-2 becomes the active CGN node for the pool “4-to-4”. This is shown using the following command, where the Active field shows true. Conversely, the same output at CGN-1 shows the Active field as false.

```
*A:CGN-2# show router 1000 nat pool "4-to-4"

===============================================================================
NAT Pool 4-to-4
===============================================================================
Description                          : (Not Specified)
ISA NAT Group                        : 1
Pool type                            : largeScale
Applications                         : (None)
Admin state                          : inService
Mode                                 : napt
Port forwarding dyn blocks reserved  : 0
Port forwarding range                : 1 - 1023
Port reservation                     : 128 blocks
Block usage High Watermark (%)       : 90
Block usage Low Watermark (%)        : 20
Subscriber limit per IP address      : 65535
Active                               : true
Deterministic port reservation       : (Not Specified)
Last Mgmt Change                     : 10/06/2016 11:29:41
===============================================================================

===============================================================================
NAT address ranges of pool 4-to-4
===============================================================================
Range                                                          Drain Num-blk
-------------------------------------------------------------------------------
10.1.4.1 - 10.1.4.254                                              0
-------------------------------------------------------------------------------
No. of ranges: 1
===============================================================================

===============================================================================
NAT members of pool 4-to-4 ISA NAT group 1
===============================================================================
Member                                                      Block-Usage-% Hi
-------------------------------------------------------------------------------
1                                                           < 1           N
-------------------------------------------------------------------------------
No. of members: 1
===============================================================================

===============================================================================
Dual-Homing
===============================================================================
Type                                 : Leader
Export route                         : 192.168.0.249/32
Monitor route                        : 192.168.0.248/32
Admin state                          : inService
Dual-Homing State                    : Active
===============================================================================
```

```
===============================================================================
Dual-Homing fate-share-group
===============================================================================
Router         Pool                                                  Type
-------------------------------------------------------------------------------
vprn1000       4-to-4                                                Leader
-------------------------------------------------------------------------------
No. of pools: 1
===============================================================================
```

Although the entire 10.1.4.0/24 block is allocated for NAT outside addressing, the address range shown in the preceding output does not include the network address (10.1.4.0/24) or broadcast address (10.1.4.255/24). Therefore, the address range does not include the entire /24 prefix and has to be fragmented into a number of longer prefixes, known through protocol NAT, to cover the 10.1.4.1-10.1.4.255 range. This is shown in the route table of VPRN 1000, following, and is due to the whole subnet not being defined in the address-range configuration. (If the address range was 10.1.4.0-10.1.4.255, there would be a single entry of 10.1.4.0/24 in the route table of VPRN 1000.)

```
*A:CGN-2# show router 1000 route-table 10.1.4.0/24 longer

===============================================================================
Route Table (Service: 1000)
===============================================================================
Dest Prefix[Flags]                          Type    Proto     Age        Pref
      Next Hop[Interface Name]                                  Metric
-------------------------------------------------------------------------------
10.1.4.0/24                                 Blackh* Aggr      00h01m34s  130
       Black Hole                                               0
10.1.4.1/32                                 Remote  NAT       00h01m34s  0
       NAT outside to mda 1/2                                   0
10.1.4.2/31                                 Remote  NAT       00h01m34s  0
       NAT outside to mda 1/2                                   0
10.1.4.4/30                                 Remote  NAT       00h01m34s  0
       NAT outside to mda 1/2                                   0
10.1.4.8/29                                 Remote  NAT       00h01m34s  0
       NAT outside to mda 1/2                                   0
10.1.4.16/28                                Remote  NAT       00h01m34s  0
       NAT outside to mda 1/2                                   0
10.1.4.32/27                                Remote  NAT       00h01m34s  0
       NAT outside to mda 1/2                                   0
10.1.4.64/26                                Remote  NAT       00h01m34s  0
       NAT outside to mda 1/2                                   0
10.1.4.128/26                               Remote  NAT       00h01m34s  0
       NAT outside to mda 1/2                                   0
10.1.4.192/27                               Remote  NAT       00h01m34s  0
       NAT outside to mda 1/2                                   0
10.1.4.224/28                               Remote  NAT       00h01m34s  0
       NAT outside to mda 1/2                                   0
10.1.4.240/29                               Remote  NAT       00h01m34s  0
       NAT outside to mda 1/2                                   0
10.1.4.248/30                               Remote  NAT       00h01m34s  0
       NAT outside to mda 1/2                                   0
10.1.4.252/31                               Remote  NAT       00h01m36s  0
       NAT outside to mda 1/2                                   0
```

```
10.1.4.254/32                                    Remote  NAT       00h01m36s  0
       NAT outside to mda 1/2                                          0
-------------------------------------------------------------------------------
No. of Routes: 15
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

CGN-2 advertising all of these longer prefixes to the edge router of the ISP is not wanted. The ISP may even enforce a minimum /24 prefix length. Therefore, the preceding configuration of VPRN 1000 shows an **aggregate** command for the 10.1.4.0/24 prefix with the argument **summary-only**. When at least one of the more-specific prefixes in the 10.1.4.0/24 range is populated in the route table of VPRN 1000, the aggregate becomes active and can be used by the route policy for exporting to IAR-1, while suppressing the more-specific routes.

The following output shows that outside pool prefixes are not populated in the NAT outside routing context at the standby CGN node (CGN-1). Even with the aggregate command configured at both CGN nodes, the aggregate prefix will only become active at the active CGN node. Therefore, only the active CGN node will advertise that aggregate prefix upstream.

```
*A:CGN-1# show router 1000 route-table 10.1.4.0/24 longer

===============================================================================
Route Table (Service: 1000)
===============================================================================
Dest Prefix[Flags]                               Type    Proto     Age        Pref
      Next Hop[Interface Name]                                       Metric
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No. of Routes: 0
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
```

So far, only one pool has been defined that could be used for stateless dual-homing for both NAT44 and NAT64. However, to allow for independent address management of each of these functions, a separate outside pool is created for each. The following configuration shows two pools in the **nat outside** context of VPRN 1000: the pool “4-to-4”, which is for NAT44 purposes, and the pool “6-to-4”, which is for NAT64 purposes.

The address range defined in the “6-to-4” pool is 10.1.6.1 to 10.1.6.254. As with the range 10.1.4.1 to 10.1.4.254 in the “4-to-4” pool, an aggregate command is configured to advertise the 10.1.6.0/24 prefix, while suppressing the more-specific prefixes.

With the redundancy node in the “6-to-4” pool, there are instances where traffic from a NAT inside routing context may be mapped to multiple outside pools, which in a stateless dual-homed environment may cause the NAT function to fail. For example, assume a NAT outside context has two pools, P1 and P2, where pool P1 is active and pool P2 is standby. An active pool can trigger the advertisement of inside and outside prefixes, and traffic will be attracted to this CGN node. When traffic arrives, it may be mapped to pool P1 on the active CGN node, due to the NAT mapping criteria. However, traffic may also be mapped to pool P2, due to the mapping criteria; this traffic will fail because the pool P2 is active on the redundant CGN node.

To ensure that this traffic failure does not happen, a group of pools accessed by the same inside routing context must all be active on the same CGN node simultaneously. To achieve this, SR OS uses a Pool Fate Sharing Group (PFSG). The PFSG ensures that all co-located pools accessed by the same inside routing context are either active or standby; not a combination of both. This is achieved by having a *leader* pool and *follower* pools.

If the leader pool is active, all follower pools are active. If the leader pool is standby, all follower pools are standby. This is enabled in the **redundancy** context of the “6-to-4” pool using the **follow** command. The **follow** command configures the pool as a follower and allows the user to access the routing context and outside pool of the leader pool. In the following example, pool “4-to-4” is a leader pool and pool “6-to-4” is a follower pool, which always assumes the same state as that of the leader.

```
*A:CGN-2#
configure
    service
        vprn 1000 customer 1 create
            aggregate 10.1.4.0/24 summary-only
            aggregate 10.1.6.0/24 summary-only
            nat
                outside
                    pool "4-to-4" nat-group 1 type large-scale create
                        redundancy
                            export 192.168.0.249/32
                            monitor 192.168.0.248/32
                            no shutdown
                        exit
                        mode napt
                        address-range 10.1.4.1 10.1.4.254 create
                        exit
                        no shutdown
                    exit
                    pool "6-to-4" nat-group 1 type large-scale create
                        redundancy
                            follow router 1000 pool "4-to-4"
```

```
                    exit
                    mode napt
                    address-range 10.1.6.1 10.1.6.254 create
                    exit
                    no shutdown
                exit
            exit
        exit
```

The following output shows a PFSG with leaders and followers in the operational state of pool “6-to-4” at CGN-2. The pool state is active, but the Dual-Homing Type field indicates that this pool is a follower. Therefore, the state is derived from the state of the leader pool, which is pool “4-to-4” in router 1000. The output also contains a list of all the pools that are part of the same PFSG.

```
*A:CGN-2# show router 1000 nat pool "6-to-4"

===============================================================================
NAT Pool 6-to-4
===============================================================================
Description                          : (Not Specified)
ISA NAT Group                        : 1
Pool type                            : largeScale
Applications                         : (None)
Admin state                          : inService
Mode                                 : napt
Port forwarding dyn blocks reserved  : 0
Port forwarding range                : 1 - 1023
Port reservation                     : 128 blocks
Block usage High Watermark (%)       : 90
Block usage Low Watermark (%)        : 20
Subscriber limit per IP address      : 65535
Active                               : true
Deterministic port reservation       : (Not Specified)
Last Mgmt Change                     : 10/06/2016 13:09:07
===============================================================================

===============================================================================
NAT address ranges of pool 6-to-4
===============================================================================
Range                                                          Drain Num-blk
-------------------------------------------------------------------------------
10.1.6.1 - 10.1.6.254                                              0
-------------------------------------------------------------------------------
No. of ranges: 1
===============================================================================

===============================================================================
NAT members of pool 6-to-4 ISA NAT group 1
===============================================================================
Member                                                  Block-Usage-% Hi
-------------------------------------------------------------------------------
1                                                       < 1           N
-------------------------------------------------------------------------------
No. of members: 1
===============================================================================
```

```
===============================================================================
Dual-Homing
===============================================================================
Type                                : Follower
Follow-pool                         : "4-to-4" router 1000
Dual-Homing State                   : Active
===============================================================================

===============================================================================
Dual-Homing fate-share-group
===============================================================================
Router         Pool                                                  Type
-------------------------------------------------------------------------------
vprn1000       4-to-4                                                Leader
vprn1000       6-to-4                                                Follower
-------------------------------------------------------------------------------
No. of pools: 2
===============================================================================
```

NAT Policies

NAT policies allow for definition of NAT attributes such as:

- filtering behavior (endpoint-independent or address-and-port-dependent)
- NAT mapping timeouts
- per-user session/flow limits
- configuration of Application Level Gateway (ALG) protocols
- high/low resource watermarks.

These attributes are generic NAT configuration parameters that are beyond the scope of this chapter.

A NAT policy also references the routing context and name of the outside pool used for the creation of NAT bindings associated with the policy. Therefore, if multiple outside pools are needed, multiple NAT policies must also be used. The following shows the configuration of the required NAT policies at CGN-2.

Because two outside pools exist in VPRN 1000 (the pool “4-to-4” for NAT44 and the pool “6-to-4” for NAT64), two policies are created using the **nat-policy** parameter. The **nat-policy** “NAT44” uses the **pool** keyword to access the “4-to-4” pool in **router** 1000, while the **nat-policy** “NAT64” uses the same **pool** keyword to access the “6-to-4” pool in **router** 1000. In this example, the same outside routing context is used for both NAT policies, but the outside routing contexts can also be different for each policy.

```
configure
    service
```

```
nat
    nat-policy "NAT44" create
        pool "4-to-4" router 1000
    exit
    nat-policy "NAT64" create
        pool "6-to-4" router 1000
    exit
exit
```

NAT Inside Context

The NAT inside routing context is the interface toward the customer or end user. There can be multiple NAT inside routing contexts mapped to a single NAT outside context (the relationship can be 1:1 or N:1). This is possible even if overlapping addresses are used in two or more NAT inside routing contexts because the NAT flow mapping tuple consists of the parameters {routing-instance, inside-IP, inside-port} mapped to {routing-instance, outside-IP, outside-port}.

The NAT inside routing context is responsible for two main functions:

1. Diverting some or all traffic toward the NAT function (ISA board).
2. Attracting traffic that should be subject to NAT toward it. This should be conditional because only the master CGNAT node should attract traffic toward itself.

For diverting traffic toward the NAT function, there are two approaches:

1. The first approach is to use IP filters with **action nat** to divert matched traffic into the ISA. Traffic subject to NAT can have a different inside and outside routing context, or the same routing context can be used for both inside and outside.
2. The second approach is a routing-based approach using a **destination-prefix** in the **nat inside** context. Any traffic with a destination address matching the defined destination-prefix is forwarded to the ISA for NAT. When the destination-prefix approach is used, different routing contexts must be used for inside and outside.

For NAT44, both the IP filter and destination-prefix approaches are permitted. For NAT64, the diversion to NAT is only supported using IPv6 filters. The example setup in this chapter consists of both NAT44 and NAT64. Therefore, for the purpose of standardization across the NAT44 and NAT64 functions, the IP filter-based approach is used for both.

In Figure 69, CE-2 is connected to PE-2 and is part of VPRN 200. To provide Internet access with stateless NAT dual-homing to VPRN 200, it is extended to both CGN-1 and CGN-2 as a NAT inside VRF. The following shows the configuration of VPRN 200 at CGN-2 with a similar configuration also applied at CGN-1.

Because one of the main functions of the NAT inside routing context is to attract traffic toward the (active) CGN node, this is configured in the **nat inside** context. The **redundancy** parameter provides a context for the configuration of NAT44 redundancy when the diversion to NAT is implemented with IP filters. In this **redundancy** context, the **peer** command is used to configure the address of the redundant peer (in this case CGN-1). If upstream traffic that is subject to NAT inadvertently arrives at the CGN node that is standby for the outside pool used for the NAT mapping, this parameter provides a forwarding address for that traffic. However, if destination-prefix based redirect to NAT is used instead of IP filters, only a **nat-policy** and **destination-prefix** need to be configured in the NAT inside routing context.

The **steering-route** command is optional. It allows for configuration of a (non-default) prefix/length that is only active in the routing table of the NAT inside routing context of the active CGN node. When this steering route is active in the routing table, it can either be advertised directly using the route-policy framework, or it can be used as an indirect next-hop to advertise some other prefix. This latter approach is used in the following configuration example, where the **steering-route** of 192.168.203.1/32 is used as an indirect next-hop for the **static-route-entry** of 0.0.0.0/0. This creates the following dependencies:

- If the CGN node is active, the steering route of 192.168.203.1/32 becomes active in the routing table of VPRN 200.
- When 192.168.203.1/32 is active, it becomes a valid indirect next-hop for 0.0.0.0/0, so this route also becomes active in the routing table.
- When the default route is active, it can be exported to the rest of the VPN using the route-policy framework.

In the first of three following configurations, the **vrf-export** command accesses a policy with the name **vrf200-export**. The second configuration shows the contents of that policy, where entry 10 accesses a prefix-list (vrf200-lsn44-default) containing the default route and advertises it into BGP-VPN, with the relevant Route-Target (target:64496:200) and Origin (origin:64496:200) Extended Communities attached.

The Origin Extended Community is used by the redundant CGN peer to drop the default route, as shown in entry 10 of the corresponding **vrf200-import** policy in the same configuration. The reason for dropping the default route at the standby CGN node is that the **vrf200-export** policy only requires that a default route is present in the routing table in order to source/advertise a default route itself. If the standby CGN node imports the default route from the active CGN node into the routing table, the standby will also attempt to advertise a default route, which is not wanted.

The **nat64** command provides the context to configure the NAT64 redundancy parameters. In the case of NAT64, the CGN node becomes a translator between the IPv6 and IPv4 address families and needs to advertise the NAT64 translator address that will be used by IPv6 clients to embed IPv4 addresses. In the first configuration, the address 2001:DB8:122:344::/96 is used as the NAT64 translator address. As with NAT44, the advertisement of this address is conditional and must only be advertised by the active CGN node. This is ensured because the prefix is only present in the routing table of the active CGN, not the standby CGN.

Entry 20 of the **vrf200-export** policy shown in the second configuration provides the relevant policy rules to ensure that this IPv6 prefix is advertised into BGP-VPN with the relevant Route-Target value when the prefix is present in the routing table.

The last parameter in the **nat inside** context is the **nat-policy**, which is known as the default NAT policy and must exist in the **nat inside** context. When multiple NAT policies are used in a single NAT inside routing context, the default NAT policy is used for any traffic that is not matched (using the destination-prefix for NAT44 or IPv4/IPv6 filters for NAT44/NAT64) and associated with an explicit NAT policy. The default NAT policy can reference a separate NAT policy, or it can reference a NAT policy that is already in use.

As previously described, the intention is to use IP filters to implement the diversion to NAT. The relevant IPv4 and IPv6 filter IDs (ID number 200 in both cases) are shown in the third configuration. The IPv4 filter has no match criteria in this example; it has **action nat** using **nat-policy** “NAT44”, which accesses the outside pool “4-to-4” in VPRN 1000. The IPv6 filter also has no match criteria and has **action nat**, but distinguishes between DSLite (DSLite is not supported for NAT stateless dual-homing) and NAT64, using the **nat-type** argument. The **nat-policy** that should be used is the policy “NAT64”, which accesses the outside pool “6-to-4” in VPRN 1000. These IP filters need to match traffic that ingresses the redundant CGN nodes from the MPLS side of VPRN 200 and are, therefore, applied in the **network ingress** context in the first configuration.

The remainder of the VPRN parameters are generic and are not explained here.

```
*A:CGN-2#
configure
    service
        vprn 200 customer 1 create
            vrf-import "vrf200-import"
            vrf-export "vrf200-export"
            route-distinguisher 64496:200
            auto-bind-tunnel
                resolution any
                exit
            exit
            static-route-entry 0.0.0.0/0
                indirect 192.168.203.1
                    no shutdown
                exit
```

```
                exit
                nat
                    inside
                        nat-policy "NAT44"
                        nat64
                            prefix 2001:db8:122:344::/96
                            no shutdown
                        exit
                        redundancy
                            peer 192.0.2.8
                            steering-route 192.168.203.1/32
                        exit
                    exit
                exit
                network
                    ingress
                        filter ip 200
                        filter ipv6 200
                    exit
                exit
                no shutdown
            exit

*A:CGN-2#
configure
    router
        policy-options
            begin
            prefix-list "vrf200-lsn44-default"
                prefix 0.0.0.0/0 exact
            exit
            prefix-list "vrf200-nat64-translator"
                prefix 2001:db8:122:344::/96 exact
            exit
            community "vrf200-soo" members "origin:64496:200"
            community "vrf200-export" members "target:64496:200"
            community "vrf200-import" members "target:64496:200"
            policy-statement "vrf200-export"
                entry 10
                    from
                        prefix-list "vrf200-lsn44-default"
                    exit
                    to
                        protocol bgp-vpn
                    exit
                    action accept
                        community add "vrf200-soo" "vrf200-export"
                    exit
                exit
                entry 20
                    from
                        prefix-list "vrf200-nat64-translator"
                    exit
                    to
                        protocol bgp-vpn
                    exit
                    action accept
                        community add "vrf200-export"
                    exit
```

```
                exit
            exit
            policy-statement "vrf200-import"
                entry 10
                    from
                        community "vrf200-soo"
                    exit
                    action reject
                exit
                entry 20
                    from
                        community "vrf200-import"
                    exit
                    action accept
                    exit
                exit
            exit

*A:CGN-2#
configure
    filter
        ip-filter 200 create
            entry 10 create
                action
                    nat nat-policy "NAT44"
                exit
            exit
        exit
        ipv6-filter 200 create
            entry 10 create
                action
                    nat nat-type nat64 nat-policy "NAT64"
                exit
            exit
        exit
    exit
```

Verification of the Active CGN Node

After the configuration of the inside and NAT outside routing contexts, with the associated NAT policies, the state of the stateless redundant CGN nodes can be verified; see Figure 70.

Figure 70 **Redundancy Status**

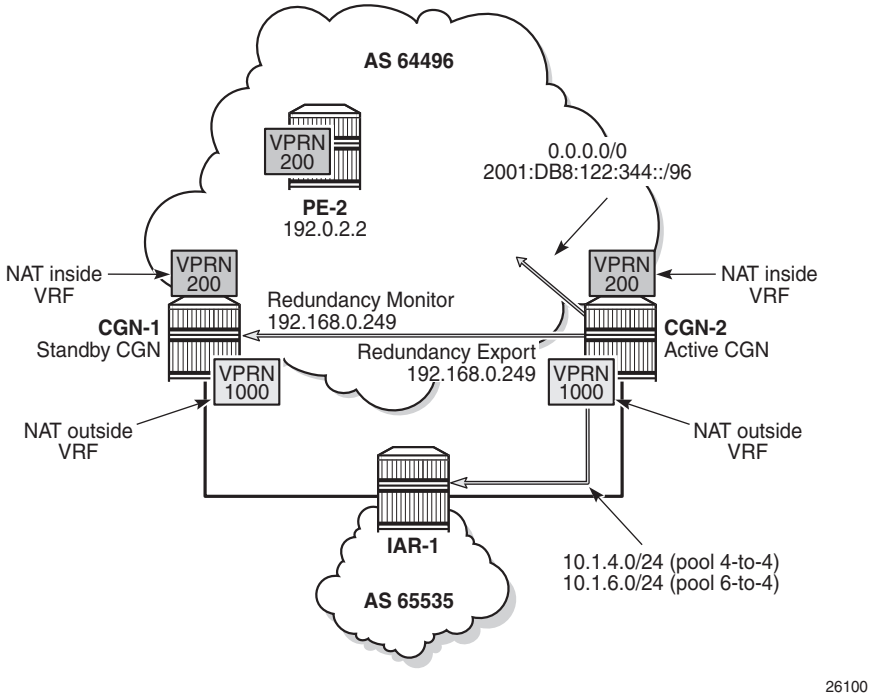

The following two outputs show that the pool “4-to-4” is a leader at both CGN-1 and CGN-2 (for which pool “6-to-4” is a follower), and that CGN-2 is the active CGN node.

```
*A:CGN-2# show router 1000 nat pool "4-to-4" | match "Dual-Homing State" pre-
lines 6
Dual-Homing
===============================================================================
Type                                  : Leader
Export route                          : 192.168.0.249/32
Monitor route                         : 192.168.0.248/32
Admin state                           : inService
Dual-Homing State                     : Active

*A:CGN-1# show router 1000 nat pool "4-to-4" | match "Dual-Homing State" pre-
lines 6
Dual-Homing
===============================================================================
Type                                  : Leader
Export route                          : 192.168.0.248/32
Monitor route                         : 192.168.0.249/32
Admin state                           : inService
Dual-Homing State                     : Standby
```

Although the redundancy export route is populated in the NAT outside routing table and advertised externally by the active CGN node (if permitted by route-policy), it is not populated in the NAT outside routing table of the standby CGN, which is not advertised externally. The following two outputs show that CGN-2 advertises its export route (192.168.0.249/32) into IPv4 BGP-VPN, but CGN-1 does not advertise its own export route because the monitor route (192.168.0.249/32) is present.

```
*A:CGN-2# show router bgp routes vpn-ipv4 rd 64496:1000 hunt
---snip---
-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : 192.168.0.249/32
Nexthop        : 192.0.2.9
Route Dist.    : 64496:1000             VPN Label      : 262128
Path Id        : None
To             : 192.0.2.10
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : 0
AIGP Metric    : None
Connector      : None
Community      : target:64496:1000
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.10
Origin         : IGP
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : N/A
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0

-------------------------------------------------------------------------------


*A:CGN-1# show router 1000 route-table 192.168.0.249/32

===============================================================================
Route Table (Service: 1000)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age         Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
192.168.0.249/32                                Remote  BGP VPN   06d03h21m  170
       192.0.2.9 (tunneled)                                          0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested


*A:CGN-1# show router bgp routes vpn-ipv4 rd 64496:1000 hunt
---snip---
-------------------------------------------------------------------------------
```

```
RIB Out Entries
------------------------------------------------------------------------------
------------------------------------------------------------------------------
```

As well as the export/monitor routes, the outside pools are populated in the NAT outside routing context and advertised by the active CGN node. The outside pools are summarized as 10.1.4.0/24 (pool “4-to-4”) and 10.1.6.0/24 (pool “6-to-4“) using the **aggregate** command in VPRN 1000. Because the (more-explicit) NAT outside pool addresses are only populated in the route table of the active CGN node, the aggregate will also only be populated in the routing table of the active CGN node. Therefore, the following policy is applied to the EBGP peering session with IAR-1 at both CGN-1 and CGN-2. The output following the policy example shows that CGN-2 advertises both of the NAT outside pools to IAR-1, while CGN-1 advertises no outside pool prefixes to IAR-1.

```
configure
    router
        policy-options
            begin
            policy-statement "vrf1000-ebgp-export"
                entry 10
                    from
                        protocol aggregate
                    exit
                    to
                        protocol bgp
                    exit
                    action accept
                        origin igp
                    exit
                exit
                default-action drop
                exit
            exit
            commit
```

```
*A:CGN-2# show router 1000 bgp neighbor 172.31.19.1 advertised-routes
===============================================================================
 BGP Router ID:192.0.2.9        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP IPv4 Routes
===============================================================================
Flag  Network                                           LocalPref   MED
      Nexthop (Router)                                  Path-Id     Label
      As-Path
-------------------------------------------------------------------------------
i     10.1.4.0/24                                       n/a         None
      172.31.19.2                                        None        -
      64496
```

```
i      10.1.6.0/24                                              n/a          None
       172.31.19.2                                               None         -
       64496
-------------------------------------------------------------------------------
Routes : 2
===============================================================================

*A:CGN-1# show router 1000 bgp neighbor 172.31.18.1 advertised-routes
===============================================================================
 BGP Router ID:192.0.2.8        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP IPv4 Routes
===============================================================================
Flag  Network                                                  LocalPref   MED
      Nexthop (Router)                                         Path-Id     Label
      As-Path
-------------------------------------------------------------------------------
No Matching Entries Found
===============================================================================
```

In the NAT inside routing instance, any routes that are used to attract traffic are populated in the relevant route tables of the active CGN node and must be advertised externally. For the NAT44 function, the steering route 192.168.203.1/32 populates the route table of the active CGN node, and this route is used as an indirect next-hop for a static-route-entry to a default route. For the NAT64 function, the NAT64 translator address 2001:DB8:122:344::/96 is used to attract IPv6 traffic with IPv4-embedded addresses.

The first of the two following outputs shows that CGN-2 is advertising the NAT44 default route as a VPN-IPv4 prefix and the NAT64 translator address as a VPN-IPv6 prefix. Both routes are advertised with the relevant Route-Target for VPRN 200 (target:64496:200) and, therefore, will be imported by PE-2 and subsequently advertised to CE-2. The second output shows the same commands entered at CGN-1 and verifies that because CGN-1 is the standby CGN node, it is not advertising either VPN-IPv4/VPN-IPv6 prefix.

```
*A:CGN-2# show router bgp routes vpn-ipv4 rd 64496:200 hunt
===============================================================================
 BGP Router ID:192.0.2.9        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
```

```
---snip---
-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : 0.0.0.0/0
Nexthop        : 192.0.2.9
Route Dist.    : 64496:200              VPN Label      : 262141
Path Id        : None
To             : 192.0.2.10
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : None
AIGP Metric    : None
Connector      : None
Community      : target:64496:200
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.10
Origin         : IGP
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : N/A
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0

-------------------------------------------------------------------------------


*A:CGN-2# show router bgp routes vpn-ipv6 rd 64496:200 hunt
===============================================================================
 BGP Router ID:192.0.2.9        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv6 Routes
===============================================================================
---snip---
-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
Network        : 2001:db8:122:344::/96
Nexthop        : ::ffff:192.0.2.9
Route Dist.    : 64496:200              VPN Label      : 262141
Path Id        : None
To             : 192.0.2.10
Res. Nexthop   : n/a
Local Pref.    : 100                    Interface Name : NotAvailable
Aggregator AS  : None                   Aggregator     : None
Atomic Aggr.   : Not Atomic             MED            : 0
AIGP Metric    : None
Connector      : None
Community      : target:64496:200
Cluster        : No Cluster Members
Originator Id  : None                   Peer Router Id : 192.0.2.10
Origin         : IGP
```

```
AS-Path        : No As-Path
Route Tag      : 0
Neighbor-AS    : N/A
Orig Validation: N/A
Source Class   : 0                      Dest Class     : 0

-------------------------------------------------------------------------------


*A:CGN-1# show router bgp routes vpn-ipv4 rd 64496:200 hunt
===============================================================================
 BGP Router ID:192.0.2.8        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv4 Routes
===============================================================================
---snip---
-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------


*A:CGN-1# show router bgp routes vpn-ipv6 rd 64496:200 hunt
===============================================================================
 BGP Router ID:192.0.2.8        AS:64496       Local AS:64496
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
                 l - leaked, x - stale, > - best, b - backup, p - purge
 Origin codes  : i - IGP, e - EGP, ? - incomplete

===============================================================================
BGP VPN-IPv6 Routes
===============================================================================
---snip---
-------------------------------------------------------------------------------
RIB Out Entries
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
```

Verification of Data Path

The host connected to CE-2 and the web server accessible from IAR-1 are used to verify the end-to-end data path for both NAT44 and NAT64.

NAT44

The host connected to CE-2 initiates an IPv4 UDP session toward the web server with a source address of 172.31.102.3 and a destination address of 10.1.128.43. The source port used is 1357, and the destination port is 80.

A two-way data transfer is verified as successful. The following output shows the details of the NAT44 binding at CGN-2, the active CGN node. The inside IP address and port are as described, while the allocated outside IP address is 10.1.4.254 using outside port 1047.

```
*A:CGN-2# tools dump nat sessions inside-ip 172.31.102.3

===============================================================================
Matched 1 session on Slot #1 MDA #2
===============================================================================
Owner               : LSN-Host@172.31.102.3
Router              : 200
Policy              : NAT44
FlowType            : UDP               Timeout (sec)       : 300
Inside IP Addr      : 172.31.102.3
Inside Port         : 1357
Outside IP Addr     : 10.1.4.254
Outside Port        : 1047
Foreign IP Addr     : 10.1.128.43
Foreign Port        : 80
Dest IP Addr        : 10.1.128.43
Nat Group           : 1
Nat Group Member    : 1
-------------------------------------------------------------------------------
===============================================================================
```

NAT64

The host connected to CE-2 also initiates an IPv6 UDP session toward the web server with a source address of 2001:DB8:4001:200::3 and a destination address of 2001:DB8:122:344::A01:802B. The destination address represents the NAT64 translator address (2001:DB8:122:344::/96) and the embedded IPv4 address (10.1.128.43) translated into colon-hexidecimal format (A01:802B). The source port is 2468 and the destination port is 80.

A two-way data transfer is verified as successful. The following output shows the details of the NAT64 binding at CGN-2. Again, the inside IPv6 address and port are as described, while the allocated outside IP address is 10.1.6.254 using outside port 1032.

```
*A:CGN-2# tools dump nat sessions inside-ip 2001:db8:4001:200::3
```

```
===============================================================================
Matched 1 session on Slot #1 MDA #2
===============================================================================
Owner               : NAT64-Sub@2001:db8:4001:200::3
Router              : 200
Policy              : NAT64
FlowType            : UDP               Timeout (sec)       : 300
Inside IP Addr      : 2001:db8:4001:200::3
                                        Inside Port         : 2468
Outside IP Addr     : 10.1.6.254
Outside Port        : 1032
Foreign IP Addr     : 10.1.128.43
Foreign Port        : 80
Dest IP Addr        : 10.1.128.43
Nat Group           : 1
Nat Group Member    : 1
-------------------------------------------------------------------------------
===============================================================================
```

Failover

Before simulating a failover test, an IPv4 UDP session is established between the host connected to CE-2 and the web server, to ensure data-path continuity during the failure.

CGN-2 is the active CGN node; to simulate a failure of the MS-ISA board, it is placed into a shutdown state.

```
*A:CGN-2# configure card 1 mda 2 shutdown

1 2016/10/13 15:10:21.39 UTC WARNING: SNMP #2004 Base 1/2/nat-in-ip
"Interface 1/2/nat-in-ip is not operational"

2 2016/10/13 15:10:21.39 UTC MINOR: NAT #2024 Base NAT
"The state of NAT group 1 changed to out-of-service."
```

Figure 71 shows the example topology in the post-failover redundancy state.

Figure 71 **Post-Failover Redundancy State**

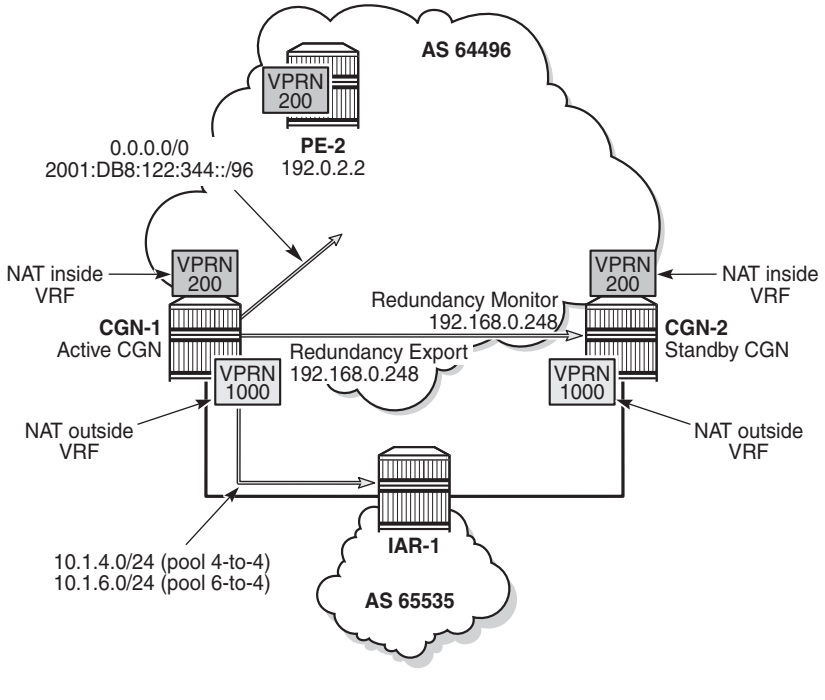

Because the MS-ISA in slot 1/2 is the only MDA in **nat-group 1**, it is sufficient to force the NAT group down. After **nat-group 1** is declared down at CGN-2, the following actions take place:

1. CGN-2 transitions the active state to false for the leader outside pool and any follower pools because its admin state changed to down, as follows:

```
22 2016/10/13 15:10:21.40 UTC WARNING: NAT #2017 vprn1000 NAT redundancy
"The Large Scale NAT activity changed to false for pool "4-to-4" - state changed
 to "Down"."

23 2016/10/13 15:10:21.40 UTC WARNING: NAT #2017 vprn1000 NAT redundancy
"The Large Scale NAT activity changed to false for pool "6-to-4" - state changed
 to "Down"."
```

2. CGN-2 withdraws the redundancy export route (192.168.0.249/32) from the NAT outside routing context (VPRN 1000). This means that the monitor route is no longer present in the routing table of CGN-1. Therefore, CGN-1 transitions to an active state and advertises its own export route. In this example, where BGP is used to advertise monitor routes, the Minimum Route Advertisement Interval is configured for 1 second, to reduce re-convergence times. However, when the monitor route is withdrawn at the standby CGN node, the system will wait for 10 seconds to ensure that this is not a route flap before declaring itself active. This is a non-configurable timer.

```
6 2016/10/13 15:10:31.53 UTC WARNING: NAT #2017 vprn1000 NAT redundancy
"The Large Scale NAT activity changed to true for pool "4-to-4" - state changed
to "Active"."

8 2016/10/13 15:10:31.53 UTC WARNING: NAT #2017 vprn1000 NAT redundancy
"The Large Scale NAT activity changed to true for pool "6-to-4" - state changed
to "Active"
```

3. Because CGN-2 is now standby, the outside pool addresses are no longer present in the routing table of the outside routing context (VPRN 1000). Therefore, they are withdrawn by CGN-2 in the EBGP peering session to IAR-1. Conversely, because CGN-1 is now active, it advertises the outside pools to IAR-1.
4. The NAT44 redundancy steering route (192.168.203.1/32) is no longer active in the NAT inside (VPRN 200) routing table at CGN-2. Therefore, the default route no longer has a valid next-hop, so the route is withdrawn. Conversely, the steering route is now present in the routing table of VPRN 200 at CGN-1. Therefore, the default route becomes active, and is advertised into VPRN 200 as a VPN-IPv4 prefix.
5. The NAT64 translator address is no longer active in the NAT inside (VPRN 200) IPv6 routing table at CGN-2, so the address is withdrawn. Conversely, the NAT64 translator address is now present in the IPv6 routing table of VPRN 200 at CGN-1 and is advertised into VPRN 200 as a VPN-IPv6 prefix.

The operational state of the outside pools can be verified at CGN-1, as follows:

```
*A:CGN-1# show router 1000 nat pool "4-to-4" | match "Dual-Homing" post-lines 12
Dual-Homing
===============================================================================
Type                                : Leader
Export route                        : 192.168.0.248/32
Monitor route                       : 192.168.0.249/32
Admin state                         : inService
Dual-Homing State                   : Active
===============================================================================

===============================================================================
Dual-Homing fate-share-group
===============================================================================
Router         Pool                                                  Type
-------------------------------------------------------------------------------
```

```
vprn1000       4-to-4                                                    Leader
vprn1000       6-to-4                                                    Follower
-------------------------------------------------------------------------------
No. of pools: 2
===============================================================================
```

Finally, the integrity of the IPv4 UDP session between the host connected to CE-2 and the web server is verified, and the associated NAT binding is shown at CGN-1, as follows:

```
*A:CGN-1# tools dump nat sessions inside-ip 172.31.102.3

===============================================================================
Matched 1 session on Slot #1 MDA #2
===============================================================================
Owner               : LSN-Host@172.31.102.3
Router              : 200
Policy              : NAT44
FlowType            : UDP               Timeout (sec)       : 300
Inside IP Addr      : 172.31.102.3
Inside Port         : 1357
Outside IP Addr     : 10.1.4.254
Outside Port        : 1049
Foreign IP Addr     : 10.1.128.43
Foreign Port        : 80
Dest IP Addr        : 10.1.128.43
Nat Group           : 1
Nat Group Member    : 1
-------------------------------------------------------------------------------
===============================================================================
```

When the failure is resolved at CGN-2 and the MS-ISA comes back online, the failover mechanism is non-revertive. This is because CGN-2 already has the CGN-1 export route present in the routing table of the NAT outside routing context (VPRN 1000) as its monitor route. The following output at CGN-2 shows the MS-ISA and NAT group 1 transitioning to in-service, followed by the active state of the outside pools changing from down to standby.

```
*A:CGN-2# configure card 1 mda 2 no shutdown


9 2016/10/13 15:55:10.39 UTC MINOR: NAT #2024 Base NAT
"The state of NAT group 1 changed to in-service."

10 2016/10/13 15:55:10.39 UTC MINOR: NAT #2025 Base NAT
"The NAT group 1 is not degraded."

17 2016/10/13 15:55:10.40 UTC WARNING: NAT #2017 vprn1000 NAT redundancy
"The Large Scale NAT activity changed to false for pool "4-to-4" - state changed
 to "Standby"."

18 2016/10/13 15:55:10.39 UTC WARNING: NAT #2017 vprn1000 NAT redundancy
"The Large Scale NAT activity changed to false for pool "6-to-4" - state changed
 to "Standby"."
```

```
19 2016/10/13 15:55:10.39 UTC MINOR: NAT #2020 Base NAT
"The NAT MDA 1/2 is now active in group 1."
```

Conclusion

NAT stateless dual-homing provides a compromise between a lack of redundancy and the protocol and state synchronization requirements for stateful NAT redundancy. This is particularly true when CGN nodes provide a gateway to the Internet where Service Level Agreements (SLAs) are often difficult to guarantee.

This chapter provides an example of how NAT stateless dual-homing is configured and describes how SR OS provides the redundancy mechanism for NAT44 and NAT64. The example in this chapter does not represent the only way that NAT stateless dual-homing can be delivered. It uses VPRNs in both the inside and outside routing contexts, but the GRT is also an option for either. It uses IP filtering for NAT diversion for both NAT44 and NAT64, but a routing-based approach using destination-prefix is also option for NAT44. It uses BGP in the NAT outside routing context and BGP-VPN in the NAT inside routing context to advertise redundancy routes externally, but any routing protocol that can be accessed through the route-policy framework is applicable.

Triple Play Service Delivery Architecture

In This Section

This section provides TPSDA configuration information for the following topics:

- RADIUS-Triggered Dynamic Data Service Provisioning
- Raw Formatting of DHCPv4/v6 Options in ESM
- Routed CO
- Subscriber Redundancy for Routed CO
- Virtual Residential Gateway Authentication Scenarios
- Virtual Residential Gateway Home Pool Management
- WiFi Aggregation and Offload — Basic Open SSID
- WiFi Aggregation and Offload — Basic Secure SSID with Distributed RADIUS Proxy
- WiFi Aggregation and Offload — IPv4/v6 Dual-Stack UEs
- WiFi Aggregation and Offload — Migrant User Support
- WiFi Aggregation and Offload — Open SSID with DSM and Lawful Intercept

ARP Hosts

This chapter describes advanced ARP host configurations.

Topics in this chapter include:

Applicability

This section describes ARP hosts and is applicable to the 7450 ESS and 7750 SR series and was tested on SR-OS 13.0 R3.

Summary

In business access area, both DHCP and PPPoE are used. However, it is possible that CPE network facing interfaces are statically configured. In such cases, the first packet the network on the user side sees is ARP to the Broadband Service Aggregator (BSA) or Broadband Service Router (BSR) interface. In order to accommodate such configurations, Enhanced Subscriber Management (ESM) feature set supports the ARP host.

In practice, this means that authentication, self-provisioning and Service Level Agreement (SLA) enforcement can be triggered by reception of ARP packets.

The BRAS node will learn the IP-MAC association based on received arp-request packet and will provision subscriber-hosts based on results from RADIUS authentication, the same way this would happen through DHCP or PPPoE.

This section provides configuration and troubleshooting commands for ARP hosts. Features common to other host types and not unique to arp-hosts are not described in this note. (Not exhaustive list: RADIUS managed routes, routed subscriber with dynamic BGP peering, Wholesale/Retail, Managed SAPs configurations, ESM related host limitation mechanisms, host High-Availability, multi-chassis peer synchronization).

Knowledge of the Triple Play Service Delivery Architecture (TPSDA) concepts is assumed throughout this document.

Warning: Enhanced Subscriber Management and RADIUS authentication are mandatory for the use of ARP hosts.

Overview

ARP host is supported in bridged CO (VPLS) and routed CO (Subscriber Interface). It is triggered by the first ARP packet from the host. ARP host is also supported in a wholesale/retail context and on managed SAPs (MSAP).

The IP and MAC addresses are extracted from the ARP request. They are copied in the access-request message sent to the RADIUS server:

- RADIUS attribute [1] Username = IP address
- VSA [26][27] Client Hardware Address = MAC address

RADIUS will, on successful authentication, reply with an access-accept message on which the ESM will create the ARP host. ESM string assignment options are out of the scope for this scenario.

Figure 72 **Bridged CO and Routed CO Example**

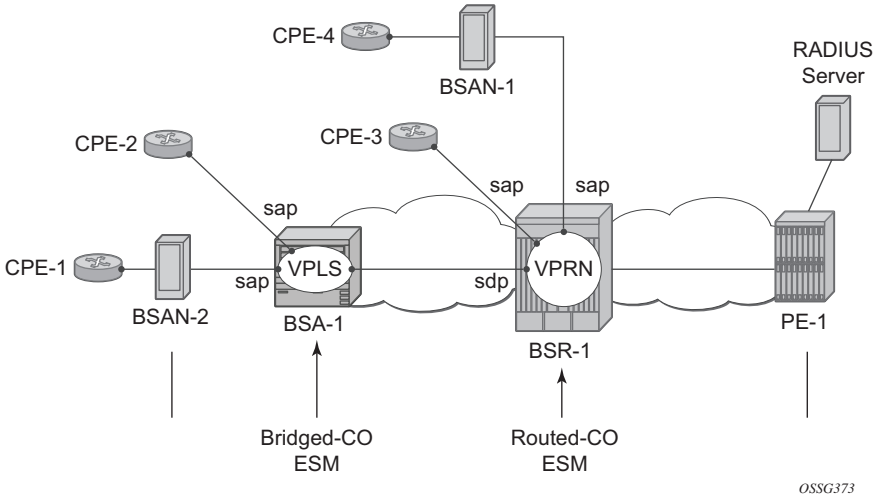

Configuration

ARP Hosts in a Bridged CO Environment

Figure 73 **ARP Hosts in a Bridged CO Environment Example**

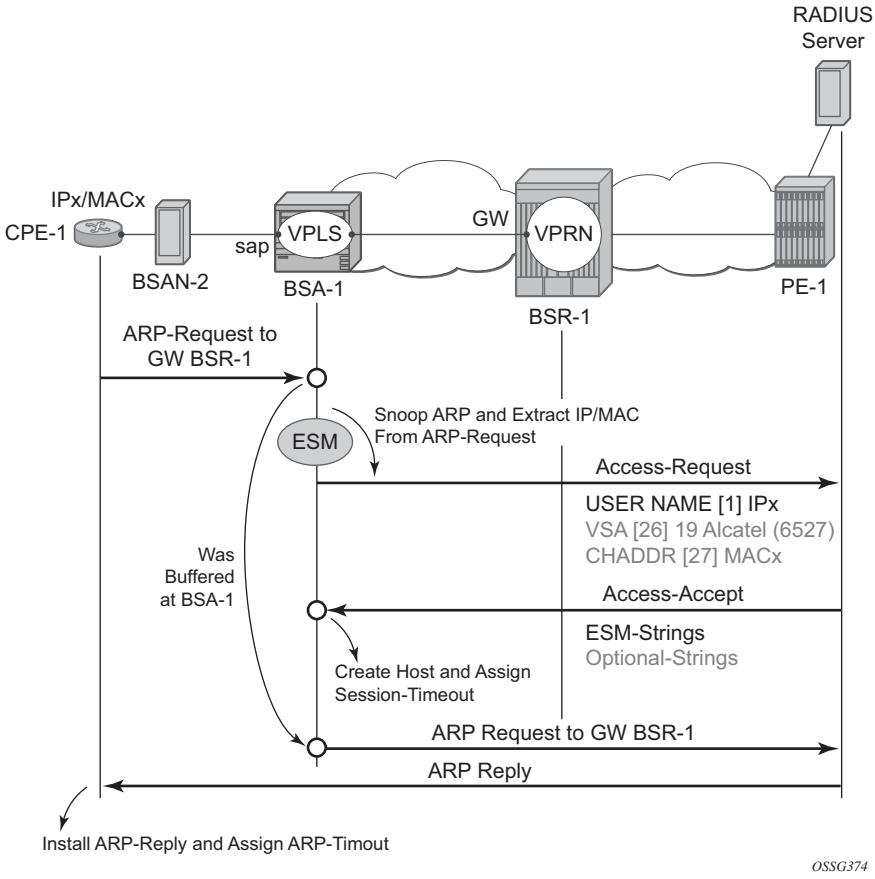

ARP-host-specific enabling for Bridged CO is achieved by a composite service; a VPLS on the BSA node and a VPRN/IES on the BSR node. RADIUS authentication and subscriber management, which mandates IP-MAC or NH-MAC type anti-spoofing, are mandatory for ARP hosts.

```
# on BSA-1
configure
    service
        vpls 2 customer 1 create
            description "ARP host - Bridged CO"
            stp
```

```
            shutdown
        exit
        sap 1/1/1:1 create
            authentication-policy "authentication-1"
            anti-spoof ip-mac
            sub-sla-mgmt
                sub-ident-policy "sub-id-default"
                multi-sub-sap 10
                no shutdown
            exit
            arp-host
                no shutdown
            exit
        exit
        spoke-sdp 12:2 create
        exit
        no shutdown
    exit
```

The RADIUS authentication policy does not require specific parameter settings. The RADIUS username attribute will contain always the host IP address which makes the authentication policy parameter user-name-format irrelevant for ARP hosts.

```
configure
    subscriber-mgmt
           authentication-policy "authentication-1" create
            password ALU
            radius-authentication-server
                server 1 address 172.16.1.1 secret ALU
            exit
            re-authentication           # optional if re-authentication is required
            accept-authorization-
change # optional if RADIUS Disconnect is required
    exit
```

The CPE ARPs are snooped and the first CPE ARP triggers a RADIUS accept-request and subsequent ARPs will trigger RADIUS re-authentication only if the ARP host configurable min-auth-interval is expired and the above re-authentication parameter is set. The initial ARP is only forwarded to the BSR-1 upon successful RADIUS authentication by means of a RADIUS access-accept message. The same RADIUS access-accept message and passing the several session limit checks, triggers the creation of the host.

The BSR-1 node requires a VPRN/IES as part of the composite service. No ARP-host-specific parameters are required on the BSR-1 for the bridged CO model.

```
# on BSR-1
configure service
    vprn 1 customer 1 create
        route-distinguisher 64496:1
        auto-bind-tunnel
```

```
        resolution-filter
            ldp
        exit
        resolution filter
    exit
    vrf-target target:64496
    interface "int-BSA1-p2mp-1" create
        description "ARP host - Bridged CO" address 10.2.0.6/29
        ip-mtu 1500
        spoke-sdp 21:2 create
        exit
    exit
```

ARP Hosts in a Routed CO Environment

Figure 74 **ARP Hosts in a Routed CO Environment Example**

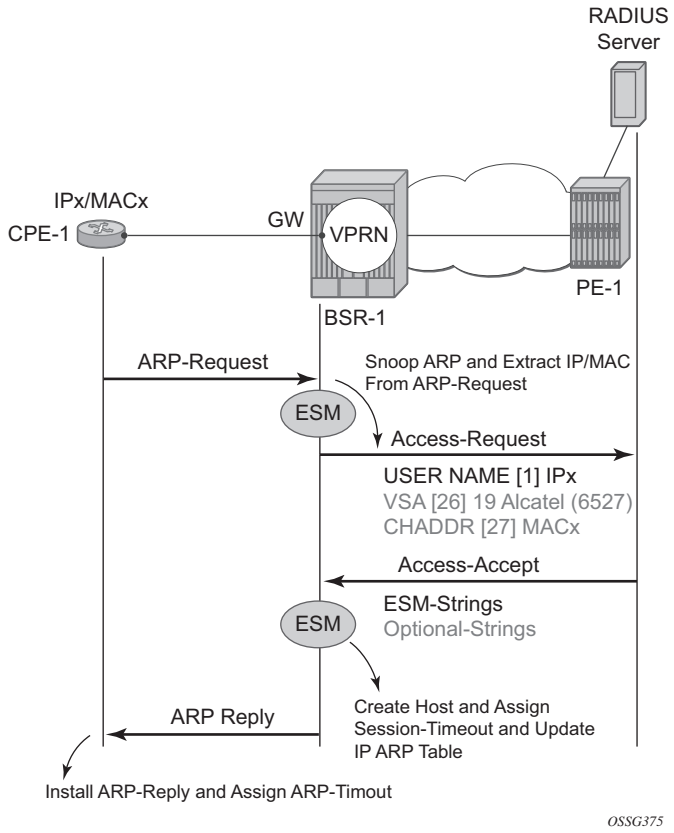

ARP-host-specific enabling for routed CO is identical for VPRN and IES services. RADIUS authentication and subscriber management, which mandates IP-MAC or NH-MAC type anti-spoofing, are mandatory for ARP hosts.

The initial ARP will, only upon successful RADIUS authentication and passing the several sessions limit checks, create the ARP host. The ARP reply or update of the IP ARP table is not performed on any unsuccessful RADIUS authentication.

```
# on BSR-1
configure service
    vprn 1 customer 1 create
        route-distinguisher 64496:1
        auto-bind-tunnel
            resolution-filter
                ldp
            exit
            resolution filter
        exit

        vrf-target target:64496:1
        subscriber-interface "sub-int-1" create
            description "ARP host - Routed CO" address 10.1.0.6/29
            group-interface "group-int-1" create
                authentication-policy "authentication-1"
                sap 1/1/1:1 create
                    anti-spoof ip-mac
                    sub-sla-mgmt
                        sub-ident-policy "sub-id-default"
                        no shutdown
                    exit
                exit
                arp-host
                    no shutdown
                exit
            exit
```

RADIUS User Configuration Bridged/Routed CO

The username in the RADIUS access request is always the statically configured IP address from the CPE and configured as key in the RADIUS users file. The RADIUS Framed-Route attribute is not required and is silently ignored (if returned to BSA/BSR node).

```
"10.1.0.1"        Auth-Type := Local, User-Password == ALU
                  Alc-Subsc-ID-Str = "arp-host-routed-%{User-name}",
                  Alc-Subsc-Prof-Str = "sub-profile-1",
                  Alc-SLA-Prof-Str = "sla-profile-1"

"10.2.0.1"        Auth-Type := Local, User-Password == ALU
                  Alc-Subsc-ID-Str = "arp-host-bridged-%{User-name}",
                  Alc-Subsc-Prof-Str = "sub-profile-1",
                  Alc-SLA-Prof-Str = "sla-profile-1"
```

Setup and Debugging of ARP Host

Identical methodologies, for bridged or Routed CO, are used to debug/setup and troubleshoot ARP hosts. Routed CO is used as an example through the rest of this section on ARP hosts.

There are two modes of ARP host debugging: all and dropped-only. The dropped-only mode shows all cases where the creation of the ARP host will be unsuccessful.

By default, all ARP hosts enabled under a service will be monitored. More specific filtering on a particular IP, MAC or SAP is optional.

All main traps are by default cyclically logged in log-id 99 and can be viewed anytime.

```
debug service id 1 arp-host mode all
```

ARP host mandate RADIUS authentication and a separate debug option is available for RADIUS interaction.

```
debug radius detail
```

CPE-3 with statically configured IP1 10.1.0.1 sends an ARP to the BSR-1 gateway.

```
1 2015/06/22 15:48:00.72 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.16.1.1:1812 id 2 len 79 vrid 1 pol authentication-1
    USER NAME [1] 8 10.1.0.1
    PASSWORD [2] 16 gy3yhtT5dF9YYilHtiiNNk
    NAS IP ADDRESS [4] 4 192.0.2.2
    VSA [26] 19 Alcatel(6527)
      CHADDR [27] 17 00:00:0a:01:00:01
"

2 2015/06/22 15:48:00.74 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 2 len 94 from 172.16.1.1:1812 vrid 1 pol authentication-1
    VSA [26] 26 Alcatel(6527)
      SUBSC ID STR [11] 24 arp-host-routed-10.1.0.1
    VSA [26] 15 Alcatel(6527)
      SUBSC PROF STR [12] 13 sub-profile-1
    VSA [26] 15 Alcatel(6527)
      SLA PROF STR [13] 13 sla-profile-1
"

3 2015/06/22 15:48:00.75 CEST MINOR: DEBUG #2001 vprn1 ARP Host
"ARP Host: Created ARP host
   VPRN 1, SAP 1/1/1:1

   IP: 10.1.0.1
   MAC: 00:00:0a:01:00:01
"

A:BSR-1# show log log-id 99
```

```
--- snipped ---
58 2015/06/22 15:48:00.73 CEST WARNING: SVCMGR #2500 Base Subscriber created
"Subscriber arp-host-routed-10.1.0.1 has been created in the system"
```

The user name in the RADIUS access-request contains the CPE's IP address independent from the user-name-format defined in the authentication policy. The MAC address of the ARP host is included in the RADIUS access-request as VSA (Alc-Client-Hardware-Addr) independent on the include-radius-attribute mac-address parameter from the authentication policy.

The **show service id 1 arp-host** command displays all active ARP hosts on this service.

```
A:BSR-1# show service id 1 arp-host
===============================================================================
ARP host table, service 1
===============================================================================
IP Address      Mac Address       Sap Id              Remaining           MC
                                                      Time                Stdby
-------------------------------------------------------------------------------
10.1.0.1        00:00:0a:01:00:01 1/1/1:1             03h59m23s
-------------------------------------------------------------------------------
Number of ARP hosts : 1
===============================================================================
A:BSR-1#
```

More specific filters such as **sap**, **ip-address**, **mac** and others can be used to show dedicated ARP hosts created on the BSR.

```
A:BSR-1# show service id 1 arp-host ip-address 10.1.0.1 detail
===============================================================================
ARP hosts for service 1
===============================================================================
Service ID           : 1
IP Address           : 10.1.0.1
MAC Address          : 00:00:0a:01:00:01
Subscriber-interface : sub-int-1
Group-interface      : group-int-1
SAP                  : 1/1/1:1
Remaining Time       : 03h59m15s

Sub-Ident            : "arp-host-routed-10.1.0.1"
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-1"
App-Profile-String   : ""
ARP host ANCP-String : ""
ARP host Int Dest Id : ""
Category-Map-Name    : ""

RADIUS-User-Name     : "10.1.0.1"

Session Timeout (s)  : 14400
Start Time           : 06/22/2015 15:48:00
```

```
Last Auth            : 06/22/2015 15:48:00
Last Refresh         : 06/22/2015 15:48:00
Persistence Key      : N/A
-------------------------------------------------------------------------------
Number of ARP hosts : 1
===============================================================================
A:BSR-1#
```

Dynamically created ARP hosts are added as /32 addresses in the routing table marked with protocol type Sub Mgmt. Routes with this protocol type are not exported into vpn-ipv4 by the default vrf-target policy. A separate vrf-export policy is required to achieve this.

```
A:BSR-1# show router 1 route-table 10.1.0.0/24 longer
===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age         Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.1.0.0/29                                     Local   Local     00h04m21s  0
       sub-int-1                                                       0
10.1.0.1/32                                     Remote  Sub Mgmt  00h00m56s  0
       [group-int-1]                                                   0
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
A:BSR-1#
```

Specific ARP host counters can be shown or cleared using the CLI command **show/clear service id 1 ARP host statistics**.

```
A:BSR-1# show service id 1 arp-host statistics
===============================================================================
ARP host statistics
===============================================================================
Num Active Hosts             : 1
Received Triggers            : 5
Ignored Triggers             : 3
Ignored Triggers (overload)  : 0
SHCV Checks Forced           : 0
Hosts Created                : 1
Hosts Updated                : 1
Hosts Deleted                : 0
Authentication Requests Sent : 4
===============================================================================
A:BSR-1#
```

The ARP hosts mandate Enhanced Subscriber managed (ESM) and therefore an anti-spoofing configuration (IP-MAC or NH-MAC). The anti-spoofing table with active hosts can be viewed with the command **show service id 1 subscriber-hosts**.

```
A:BSR-1# show service id 1 subscriber-hosts
===============================================================
Subscriber Host table
===============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID Origin       Fwding State
---------------------------------------------------------------
1/1/1:1                arp-host-routed-10.1.0.1
  10.1.0.1
    00:00:0a:01:00:01    N/A       ARP-Host     Fwding
---------------------------------------------------------------
Number of subscriber hosts : 1
===============================================================
A:BSR-1#
```

An ARP host can be manually deleted from the system using one of the two following methods:

- clear service id 1 arp-host
- RADIUS disconnect message

Using the first method, **clear service id 1 arp-host** and omitting any more specific parameter than ARP host will result in the removal of all ARP hosts in this service. Extra filters like **ip-address**, **mac** or **sap-id** are used to remove a specific ARP host.

```
*A:BSR-1# clear service id 1 arp-host
  - arp-host {all | mac <ieee-address> | sap <sap-id> | ip-address <ip-address[/
mask]> }
  - arp-host {port <port-id> | {inter-dest-id <intermediate-destination-id> | no-
inter-dest-id} [port <port-id>] }
  - arp-host statistics [sap <sap-id> | interface <interface-name>]

A:BSR-1# clear service id 1 arp-host ip-address 10.1.0.1
```

Using the second method, RADIUS disconnect always result in the removal of a unique host because **nas-port-id** and **framed-ip-address** are mandatory parameters in the RADIUS disconnect message. This RADIUS disconnect message is used also for other host-types.

```
nas-port-id = 1/1/1:1
framed-ip-address=10.1.0.1
```

RADIUS disconnect messages are, for security reasons, rejected by default and are allowed instead of enabled by setting **accept-authorization-change** parameter in the authentication policy. The **debug radius detail** command and **show subscriber-mgmt authentication coa-statistics** can be used during troubleshooting.

```
10 2015/06/22 15:51:08.43 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Disconnect Request(40) id 247 len 44 from 172.16.1.1:46749 vrid 1
    SESSION ID [44] 22 02DAFF0000000255881288
"

11 2015/06/22 15:51:08.42 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Disconnect Ack(41) 172.16.1.1:46749 id 247 len 26 vrid 1 pol authentication-1
    TERMINATE CAUSE [49] 4 Admin Reset(6)
"

12 2015/06/22 15:51:08.43 CEST MINOR: DEBUG #2001 vprn1 ARP Host
"ARP Host: Removed ARP host
   VPRN 1, SAP 1/1/1:1

   IP: 10.1.0.1
   MAC: 00:00:0a:01:00:01
"
```

In both cases the ARP host with an IP address is removed from the system together with all related state information (such as an anti-spoof filter and an IP ARP entry).

ARP Host Session Timeout

The ARP host session timeout is a time value between 300 and 14400 seconds and becomes the remaining time at the moment the first ARP request results in a successful host creation.

The host is removed from the system at the moment the remaining time becomes zero. The reset of remaining time to session timeout is done by any subsequent arp-request or arp-reply for this host.

The default assigned session timeout at ARP host creation time is 14400 seconds but this value can be overruled by the optional RADIUS attribute session-Timeout and not by the node group-interface arp-timeout parameter.

RADIUS values lower than 300 seconds will be silently adjusted to 300 seconds and values above 14400 seconds are topped silently to 14400 seconds.

```
"10.1.0.1"    Auth-Type := Local, User-Password == ALU
              Alc-Subsc-ID-Str = "arp-host-routed-%{User-name}",
```

```
                Alc-Subsc-Prof-Str = "sub-profile-1",
                Alc-SLA-Prof-Str = "sla-profile-1",
                Session-Timeout = 300  # value in seconds


A:BSR-1# show service id 1 arp-host
===============================================================================
ARP host table, service 1
===============================================================================
IP Address      Mac Address       Sap Id              Remaining           MC
                                                      Time                Stdby
-------------------------------------------------------------------------------
10.1.0.1        00:00:0a:01:00:01 1/1/1:1             00h04m56s
-------------------------------------------------------------------------------
Number of ARP hosts : 1
===============================================================================
A:BSR-1#

A:BSR-1# show service id 1 arp-host ip-address 10.1.0.1 detail
===============================================================================
ARP hosts for service 1
===============================================================================
Service ID           : 1
IP Address           : 10.1.0.1
MAC Address          : 00:00:0a:01:00:01
Subscriber-interface : sub-int-1
Group-interface      : group-int-1
SAP                  : 1/1/1:1
Remaining Time       : 00h04m39s

Sub-Ident            : "arp-host-routed-10.1.0.1"
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-1"
App-Profile-String   : ""
ARP host ANCP-String : ""
ARP host Int Dest Id : ""
Category-Map-Name    : ""

RADIUS-User-Name     : "10.1.0.1"

Session Timeout (s)  : 300
Start Time           : 06/22/2015 15:53:11
Last Auth            : 06/22/2015 15:53:11
Last Refresh         : 06/22/2015 15:53:11
Persistence Key      : N/A
-------------------------------------------------------------------------------
Number of ARP hosts : 1
===============================================================================
A:BSR-1#
```

Typical time related parameters of the ARP host are:

Table 9 **ARP Host Time-Related Parameters**

Parameter	Comment
Session Timeout	Time value in seconds and retrieved by default or by the RADIUS Accept message and pasted into the remaining time at the moment of ARP host creation or RADIUS re-authentication.
Remaining Time	The remaining time before the ARP host is deleted from the system (updated each time an ARP request/reply is seen for this host).
Start Time	Time and date when this host was created (first ARP received).
Last Auth	Time and date when this host was last RADIUS authenticated.
Last Refresh	Time and date when last ARP was received for this host.

ARP hosts do not have an expiry timer in the ARP table and have type **managed**.

```
A:BSR-1# show service id 1 arp 10.1.0.1
===============================================================================
ARP Table
===============================================================================
IP Address      MAC Address       Type    Expiry    Interface         SAP
-------------------------------------------------------------------------------
10.1.0.1        00:00:0a:01:00:01 Managed 00h00m00s group-int-1       1/1/1:1
===============================================================================
A:BSR-1#
```

An automatic mechanism is foreseen to handle the possible asynchrony between the ARP session timeout values installed on the BSR and the ARP timeouts installed on the CPE. This mechanism is mostly effective in case the timeout on the CPE exceeds the timeout on the BSR. In this case, the BSR session would expire, resulting in a host removal with a deletion of the corresponding anti-spoof entry because the CPE ARP request arrives too late. This CPE ARP request will however recreate the session but requires the complete setup of the host RADIUS authentication included. This mechanism causes unwanted service interruption for this ARP host.

A better approach, which is implemented in an automatic way, and illustrated in Figure 75 is an ARP request triggered from the BSR towards the CPE prior to the session timeout. A CPE ARP reply will then reset the remaining lifetime of the ARP host to the session timeout. If the ARP reply is received outside the **min-auth-interval** window and the parameter re-authentication from the authentication policy is set, than RADIUS re-authentication is executed. This re-authentication mechanism is described further in the throttling toward the RADIUS section.

Figure 75 **ARP Host Session Timeout Example**

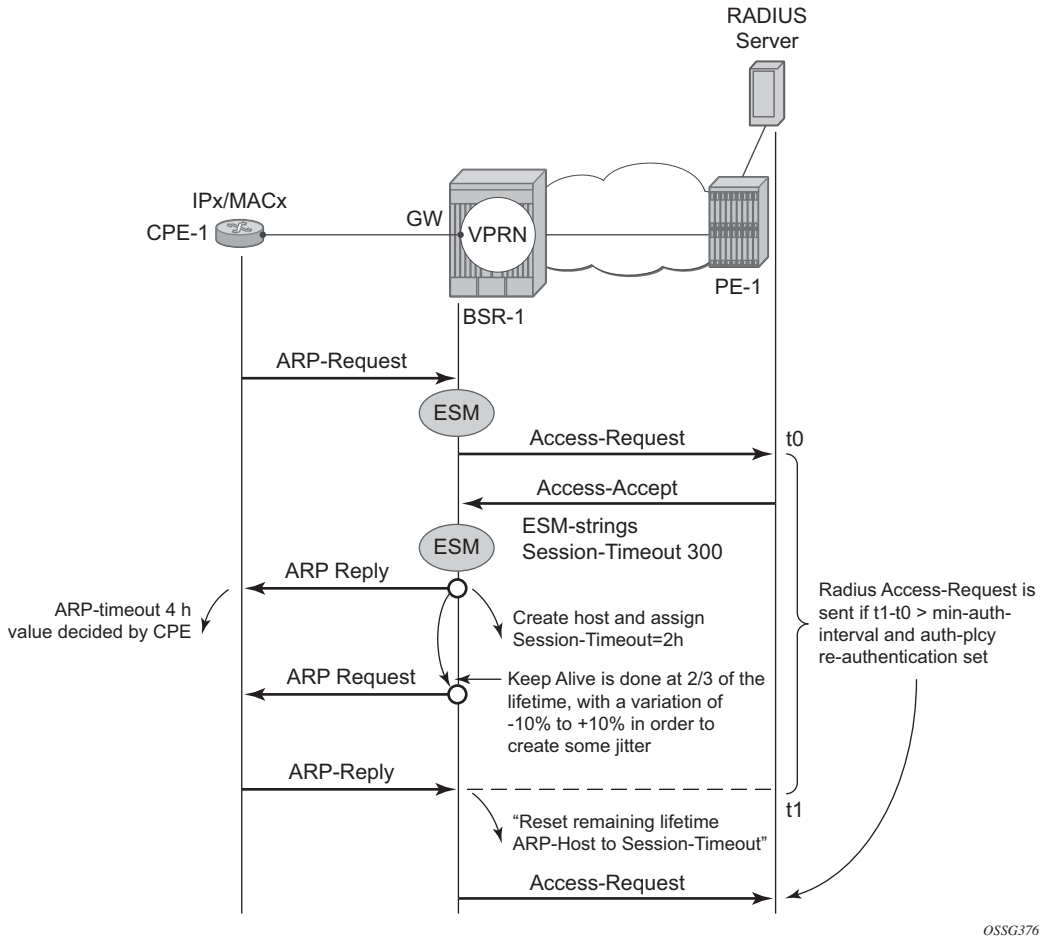

This mechanism, also known as automatic Subscriber Host Connectivity Verification (SHCV), will prevent that the host will be deleted and re-created, resulting in undesired service interruptions, in case asynchronous CPE-BSR ARP session values would be used.

The **debug service id 1 host-connectivity-verify** command shows the sequence of events and can be used during troubleshooting. Debugging and ARP host counters show the automatic SHCV mechanism with an active CPE.

```
4 2015/06/25 16:32:45.21 CEST MINOR: DEBUG #2001 vprn1 SHCV
"SHCV: Forced Check Scheduled
   1/1/1:1
   ARP host 10.1.0.1 00:00:0a:01:00:01"

5 2015/06/25 16:32:46.11 CEST MINOR: DEBUG #2001 vprn1 SHCV
"SHCV: Forced Check
   1/1/1:1
```

```
   ARP host 10.1.0.1 00:00:0a:01:00:01"

6 2015/06/25 16:32:46.12 CEST MINOR: DEBUG #2001 vprn1 SHCV
"SHCV: Received Reply
   1/1/1:1
   ARP host 10.1.0.1 00:00:0a:01:00:01"

7 2015/06/25 16:32:46.12 CEST MINOR: DEBUG #2001 vprn1 ARP Host
"ARP Host: Updated ARP host
   VPRN 1, SAP 1/1/1:1

   IP: 10.1.0.1
   MAC: 00:00:0a:01:00:01
"

8 2015/06/25 16:34:29.34 CEST MINOR: DEBUG #2001 vprn1 SHCV
"SHCV: Received Reply
   1/1/1:1
   ARP host 10.1.0.1 00:00:0a:01:00:01"

9 2015/06/25 16:34:29.34 CEST MINOR: DEBUG #2001 vprn1 ARP Host
"ARP Host: Updated ARP host
   VPRN 1, SAP 1/1/1:1

   IP: 10.1.0.1
   MAC: 00:00:0a:01:00:01
"


A:BSR-1# show service id 1 arp-host statistics
===============================================================================
ARP host statistics
===============================================================================
Num Active Hosts             : 1
Received Triggers            : 3
Ignored Triggers             : 0
Ignored Triggers (overload)  : 0
SHCV Checks Forced           : 1
Hosts Created                : 1
Hosts Updated                : 2
Hosts Deleted                : 0
Authentication Requests Sent : 1
===============================================================================
A:BSR-1#
```

CPEs that are not active and therefore do not respond to ARP requests as part of the automatic SHCV check will be rechecked three times with 10 second intervals.

The number of retries and the interval cannot be changed. A trap is generated, but the ARP host is not removed and will remain until the session-timeout expires or until the host will revive. This mechanism is displayed in Figure 76.

Figure 76 **Trap Generation Example**

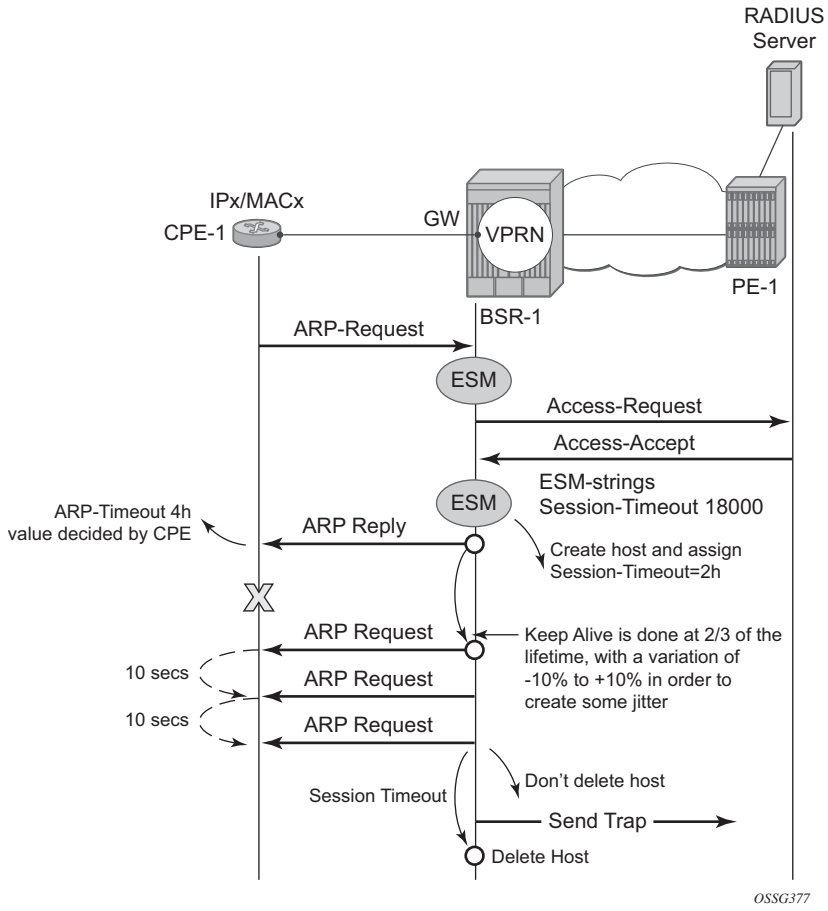


```
16 2015/06/25 16:42:38.21 CEST MINOR: DEBUG #2001 vprn1 SHCV
"SHCV: Forced Check Scheduled
   1/1/1:1
   ARP host 10.1.0.1 00:00:0a:01:00:01"

17 2015/06/25 16:42:39.11 CEST MINOR: DEBUG #2001 vprn1 SHCV
"SHCV: Forced Check
   1/1/1:1
   ARP host 10.1.0.1 00:00:0a:01:00:01"

18 2015/06/25 16:42:49.11 CEST MINOR: DEBUG #2001 vprn1 SHCV
"SHCV: Forced Check
   1/1/1:1
   ARP host 10.1.0.1 00:00:0a:01:00:01"

19 2015/06/25 16:42:59.11 CEST MINOR: DEBUG #2001 vprn1 SHCV
"SHCV: Forced Check
   1/1/1:1
   ARP host 10.1.0.1 00:00:0a:01:00:01"

20 2015/06/25 16:43:09.11 CEST MINOR: DEBUG #2001 vprn1 SHCV
```

```
"SHCV: Connectivity Lost
   1/1/1:1
   ARP host 10.1.0.1 00:00:0a:01:00:01"

21 2015/06/25 16:43:10.21 CEST MINOR: DEBUG #2001 vprn1 ARP Host
"ARP Host: Removed ARP host
   VPRN 1, SAP 1/1/1:1

   IP: 10.1.0.1
   MAC: 00:00:0a:01:00:01
"
```

Throttling Toward RADIUS

A new ARP request from the ARP host will trigger RADIUS re-authentication only when the min-auth-interval is expired. The minimum RADIUS authentication interval between two consecutive authentication attempts for the same ARP host is by default 15 minutes but can range between 1 and 6000 minutes.

```
configure
    service
       vprn 1 customer 1 create
         --snip--
                arp-host
                      min-auth-interval 60   # value in minutes
                      no shutdown
                  exit
                exit
```

Figure 77 **Throttling Toward RADIUS Example**

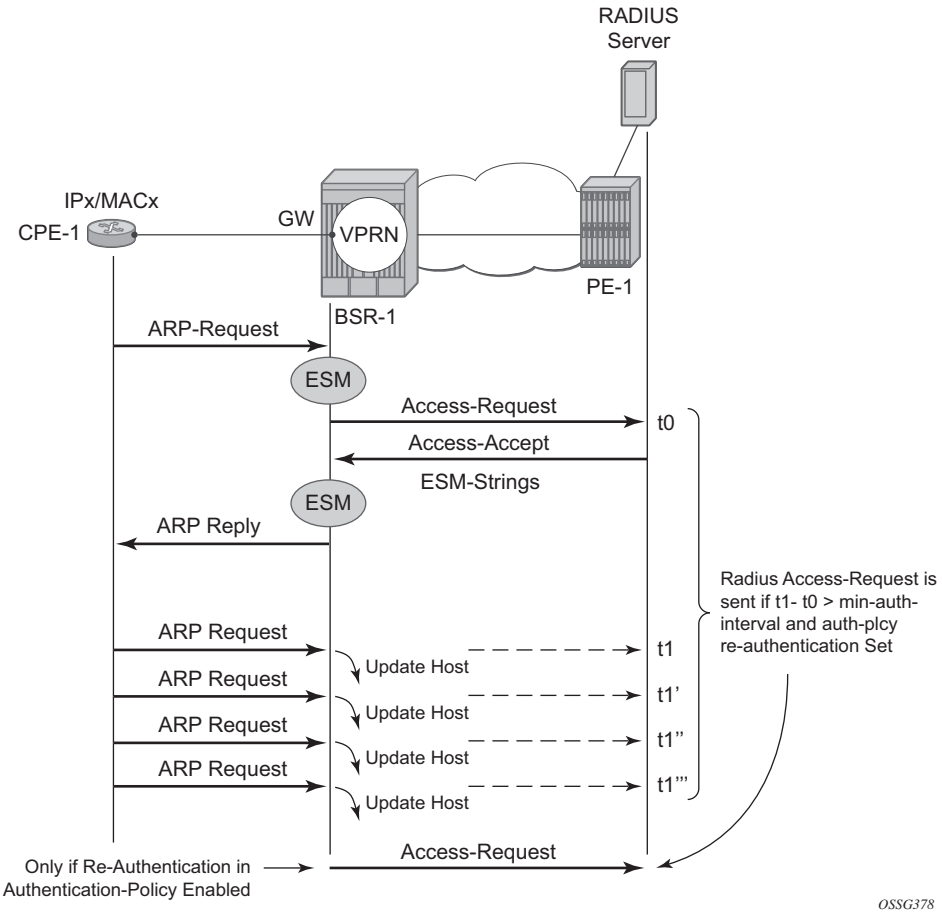


```
A:BSR-1# show service id 1 arp-host detail
===============================================================================
ARP hosts for service 1
===============================================================================
Service ID           : 1
IP Address           : 10.1.0.1
MAC Address          : 00:00:0a:01:00:01
Subscriber-interface : sub-int-1
Group-interface      : group-int-1
SAP                  : 1/1/1:1
Remaining Time       : 00h04m31s

Sub-Ident            : "arp-host-routed-10.1.0.1"
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-1"
App-Profile-String   : ""
ARP host ANCP-String : ""
ARP host Int Dest Id : ""
Category-Map-Name    : ""
```

```
RADIUS-User-Name      : "10.1.0.1"

Session Timeout (s)   : 300
Start Time            : 06/22/2015 15:59:32
Last Auth             : 06/22/2015 15:59:32
Last Refresh          : 06/22/2015 16:00:33
Persistence Key       : N/A
-------------------------------------------------------------------------------
Number of ARP hosts : 1
===============================================================================
A:BSR-1#
```

ARP Host Mobility

In order for ARP host mobility to function, host-connectivity-verification must **not** be enabled. This is different compared to DHCP host mobility.

The implementation for routed CO is displayed in Figure 78 and works the same for bridged CO. The **mac-pinning** command in routed CO context has no influence on this behavior.

Figure 78 **ARP Host Mobility Example**

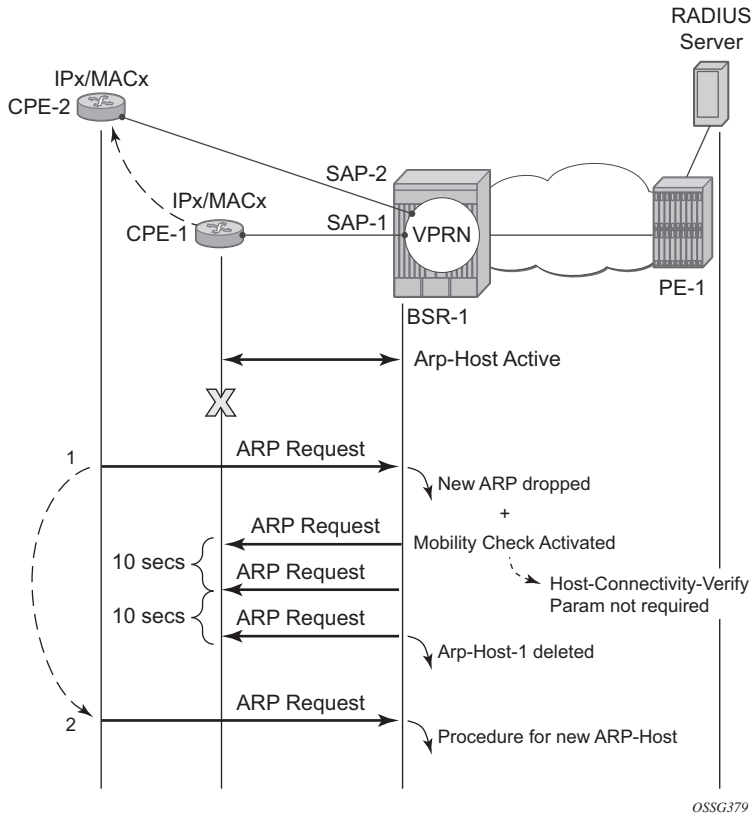

ARP Host Persistency

ARP hosts can be made persistent across reboots and do not differ with other host types like DHCP hosts.

```
configure system
        persistence
            subscriber-mgmt
                location cf3:
            exit
        exit
```

The persistence key and the index into the persistency file are linked to the ARP host at host creation time.

```
A:BSR-1# show service id 1 arp-host detail
===============================================================================
ARP hosts for service 1
===============================================================================
Service ID           : 1
IP Address           : 10.1.0.1
MAC Address          : 00:00:0a:01:00:01
Subscriber-interface : sub-int-1
Group-interface      : group-int-1
SAP                  : 1/1/1:1
Remaining Time       : 00h04m32s

Sub-Ident            : "arp-host-routed-10.1.0.1"
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-1"
App-Profile-String   : ""
ARP host ANCP-String : ""
ARP host Int Dest Id : ""
Category-Map-Name    : ""

RADIUS-User-Name     : "10.1.0.1"

Session Timeout (s)  : 300
Start Time           : 06/22/2015 15:59:32
Last Auth            : 06/22/2015 15:59:32
Last Refresh         : 06/22/2015 16:03:33
Persistence Key      : 0x00000000
-------------------------------------------------------------------------------
Number of ARP hosts : 1
===============================================================================
*A:BSR-1#
```

The content of the stored record is viewed with the **tools dump persistency** command using the **persistency** key as a record number.

```
*A:BSR-1# tools dump persistence submgt record 0x00000000
----------------------------------------
Persistence Record
----------------------------------------
```

```
Client      : submgt
Persist-Key : 0x00000000
Filename    : cf3:\submgmt.011
Entries     : Index  FedHandle  Last Update               Action Valid
              000040 0x00000000 2015/06/22 14:01:31 (UTC) ADD    Yes
Data        : 243 bytes

 Host Type      : ARP host
 Service ID     : 1
 SAP ID         : 1/1/1:1
 NH MAC         : 00:00:0a:01:00:01
 Created        : 2015/06/22 13:59:32 (UTC)
 IP             : 10.1.0.1
 Session Timeout: 300 (seconds)
 RADIUS Fallback: NO
 Acct-Sess-Id   : 02DAFF00000008558814C4
 Multi-Sess-Id  : 02DAFF00000009558814C4
 Class Attr     : 0 bytes
 User-Name      : "10.1.0.1"
 host is authenticated by radius: true
 Subscriber-Id  : "arp-host-routed-10.1.0.1"
 Sub-Profile-Str: "sub-profile-1"
 SLA-Profile-Str: "sla-profile-1"

*A:BSR-1#
```

Session Limitation Options

The maximum number of allowed arp-hosts in a bridged CO model can be configured by the per SAP parameter host-limit in the range of 1 to 32767.

```
configure
    service
       vpls 2
         --snip
           sap 1/1/3:1
                arp-host
                      host-limit 1            # default value 1
                      no shutdown
                   exit
                exit
                 exit
```

The maximum number of allowed arp-hosts in a routed CO model can be configured by the per group interface parameter host-limit in the range of 1 to 32767 and/or by the **sap-host-limit** parameter.

```
configure
    service
       vprn 1
         --snip--
                arp-host
                      host-limit 1            # default value
```

```
            sap-host-limit 1        # default value
            no shutdown
        exit
    exit
```

Warning: Specific ESM-related host limit mechanisms such as **sla-profile host-limit** and **sub-sla-mgmt multi-sub-sap** apply also for ARP hosts but are not further elaborated in this section.

Debugging **arp-host mode dropped-only** indicates the dropped reason and a logging trap is included in the standard log 99.

```
56 2015/06/22 16:08:35.37 CEST MINOR: DEBUG #2001 vprn1 ARP Host
"ARP Host: Dropped trigger
   VPRN 1, SAP 1/1/1:1

   Problem: Interface limit (1) of ARP hosts reached

   IP: 10.1.0.2
   MAC: 00:00:0a:01:00:02
"

*A:BSR-1# show log log-id 99
--- snipped ---
78 2015/06/22 16:08:55.52 CEST WARNING: SVCMGR #2520 vprn1 ARP Host Population Error
"ARP host table population error on SAP 1/1/1:1 in service 1 -
 Interface limit (1) of ARP hosts reached"
```

Increasing the **sap-host-limit** to 100 and the host-limit to 2000 results in the following summary:

```
*A:BSR-1# show service id 1 arp-host summary
===============================================================================
ARP host Summary, service 1
===============================================================================
Interface Name                   Used      Provided  Admin State
-------------------------------------------------------------------------------
group-int-1                      1         2000      inService
-------------------------------------------------------------------------------
Interfaces: 1                    1
-------------------------------------------------------------------------------
===============================================================================
*A:BSR-1#
```

Conclusion

This note provides configuration and troubleshooting commands for dynamic ARP hosts. ARP hosts can be instantiated in a Layer 2 bridged CO (VPLS) environment as well as in a Layer 3 Routed CO (IES/VPRN subscriber interface) context.

Bridged CO

This chapter provides information about Bridged CO model of Triple Play Service Delivery Architecture (TPSDA).

Topics in this chapter include:

Applicability

This chapter is applicable to the 7750 SR and 7450 ESS series and was tested on SR-OS 7.0 R4. Chassis mode B or higher must be used. The 7750 SR-c4 is supported from 8.0R4 and higher. This note is related only to the use of IPv4 DHCP hosts.

Summary

This chapter provides information about basic technology, network topology and configuration examples which are used in Bridged CO model of Triple Play Service Delivery Architecture (TPSDA). Regardless of aggregation technologies which are used by customers Nokia offers flexible and easy to use methodology to manage DHCP subscribers in Layer 2 domain and distribute subscriber management intelligence across multiple nodes.

Knowledge of the Triple Play Service Delivery Architecture (TPSDA) concepts is assumed throughout this chapter.

Overview

Bridged CO is a basic TPSDA model and implies that access nodes are united in one Layer 2 aggregation network and VPLS is used as a primary technology for these purposes. This fact allows the use of subscriber management functionality on BSA nodes. Bridged CO network topology is shown in Figure 79.

Figure 79 **Bridged CO Network Topology**

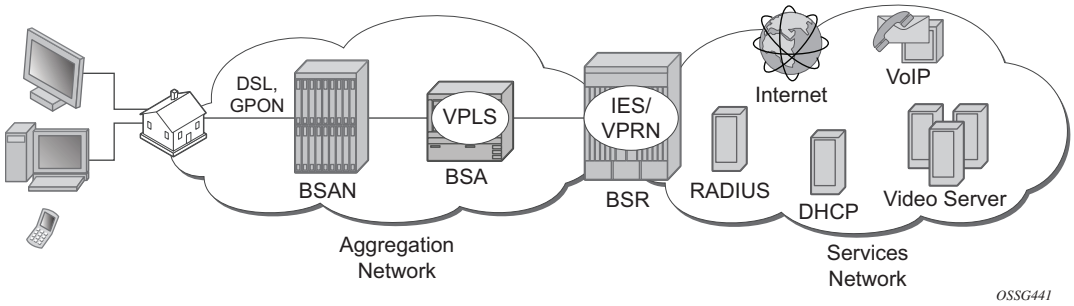

Following types of nodes are defined in Bridged CO model:

- Broadband Service Access Node (BSAN) — Access node connected to Layer 2 domain to aggregate all traffic from subscribers (IP DSLAM, ethernet switch).
- Broadband Services Aggregator (BSA) — Layer 2 node, which is capable for subscriber management in VPLS service (7450 ESS).
- Broadband Service Router (BSR) — Layer 3 node, which is capable for routing and service allocation (7750 SR).

As any model, Bridged CO introduces several key concepts that must be determined in advance. Major ones are presented in Figure 80 and include:

- Subscriber— A set of hosts belonging to a single connection line (switch port, DSL line)
- Subscriber host — Unique customer device (could be PC, IP phone, STB, routed CPE).
- Subscriber-profile — Configured entity which defines the aggregate QoS for all hosts within a subscriber context.
- SLA-profile — Configured entity which defines QoS policies and filters for a subset of hosts within a subscriber context.
- Subscriber identification policy — Configured entity which defines the python script for dynamic subscriber host identification
- Authentication policy — Configured entity which defines the RADIUS servers to use for dynamic subscriber host identification

- Subscriber identification string — 32 characters identification string which uniquely identifies a subscriber on a node.

Figure 80 **Key Concepts of Bridged CO Model**

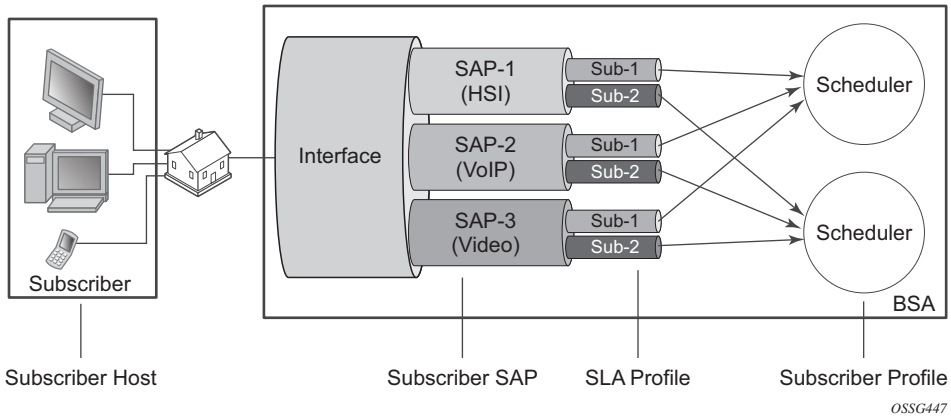

For normal operation each subscriber has to get several parameters / attributes:

- Subscriber-ID — Attribute, which uniquely identifies subscriber on the node and used as index key in subscriber database
- IP parameters — Attributes, which allows host to get access to services
- Subscriber profile and SLA — Profile for subscriber host ?a set of filters and QoS policies.
- Lease time — Period when subscriber parameters are kept in subscriber database on the node.

There are several methods how to get each of these parameters:

- Static
- Python scripts
- RADIUS
- DHCP

Each of the subscriber parameters could be defined in several ways simultaneously. In this case use the following algorithm for selecting:

Step 1. For subscriber profile

1. A lookup in t**he subscriber-explicit-map is** performed with the *sub-ident* string returned by the Python script, RADIUS or statically configured. If a matching entry is found, the sub-profile-name (if defined) is taken. If no entry was found go to 2.

```
A:BSA>config>subscr-mgmt# info
        explicit-subscriber-map
            entry key "Sub-1" sub-profile "sub-profile-1" sla-profile "sla-profile-
1"
```

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 go to 3.

```
A:BSA>config>service>vpls>sap# info
                sub-sla-mgmt
                    sub-ident-policy "sub-ident-policy-1"

A:BSA>config>subscr-mgmt# info
        sub-ident-policy "sub-ident-policy-1" create
            sub-profile-map
                entry key "sub-1" sub-profile "sub-profile-1"
```

3. If provisioned, the sub-profile-name is taken from the def-sub-profile attribute on the SAP. If no entry was found go to 4.

```
A:BSA>config>service>vpls>sap# info
                sub-sla-mgmt
                    def-sub-profile "sub-profile-1"
```

4. If a sub-profile with the name **default** is provisioned. If no entry was found DHCP Ack is dropped.

```
A:BSA>config>subscr-mgmt# info
        sub-profile "default" create
```

Figure 81 **Flow Chart for Subscriber-Profile Identification Algorithm**

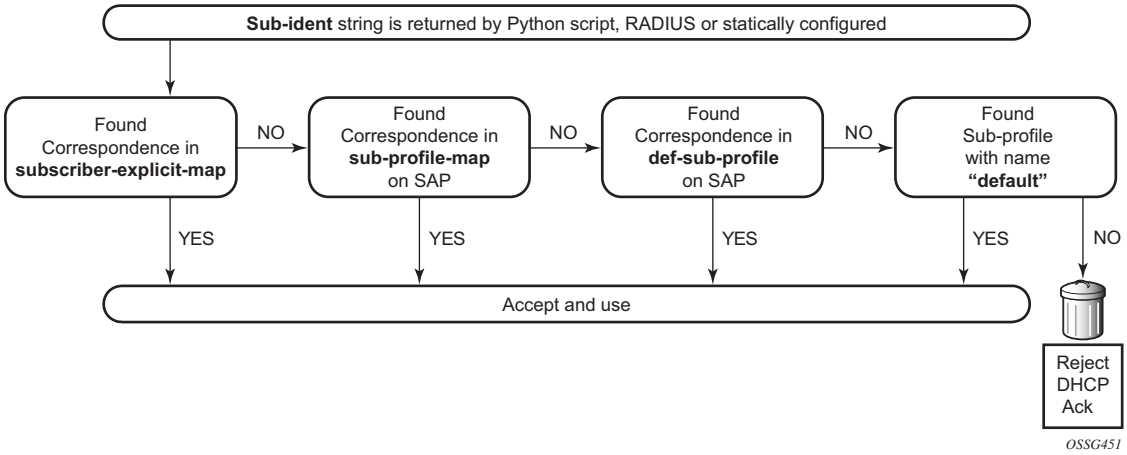

Step 2. For SLA profile

1. The sla-profile-name is taken from the sub-ident string (returned by the Python script, RADIUS or statically configured) in the subscriber-explicit-map. If no entry was found go to 2.

```
A:BSA>config>subscr-mgmt# info
        explicit-subscriber-map
            entry key "Sub-1" sub-profile "sub-profile-1" sla-profile "sla-profile-
1"
```

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 corresponding sla-profile-name is used. If no entry was found go to 3:

```
A:BSA>config>subscr-mgmt# info
        sub-profile "sub-profile-1" create
            sla-profile-map
                entry key "sla-1" sla-profile "sla-profile-1"
```

3. The sla-profile-name is taken from sla-profile-map of the sub-ident-policy configured on the SAP. The corresponding sla-profile-name is used. If no entry was found go to 4.

```
A:BSA>config>service>vpls>sap# info
                sub-sla-mgmt
                    sub-ident-policy "sub-ident-policy-1"

A:BSA>config>subscr-mgmt# info
        sub-ident-policy "sub-ident-policy-1" create
            sla-profile-map
                entry key "sla-1" sla-profile "sla-profile-1"
```

4. The *sla-profile-name* is taken from the *def-sla-profile* attribute on the SAP. If no entry was found DHCP Ack is dropped.

```
A:BSA>config>service>vpls>sap# info
                sub-sla-mgmt
                    def-sla-profile "sla-profile-1"
```

Figure 82 **Flowchart for SLA-Profile Identification Algorithm**

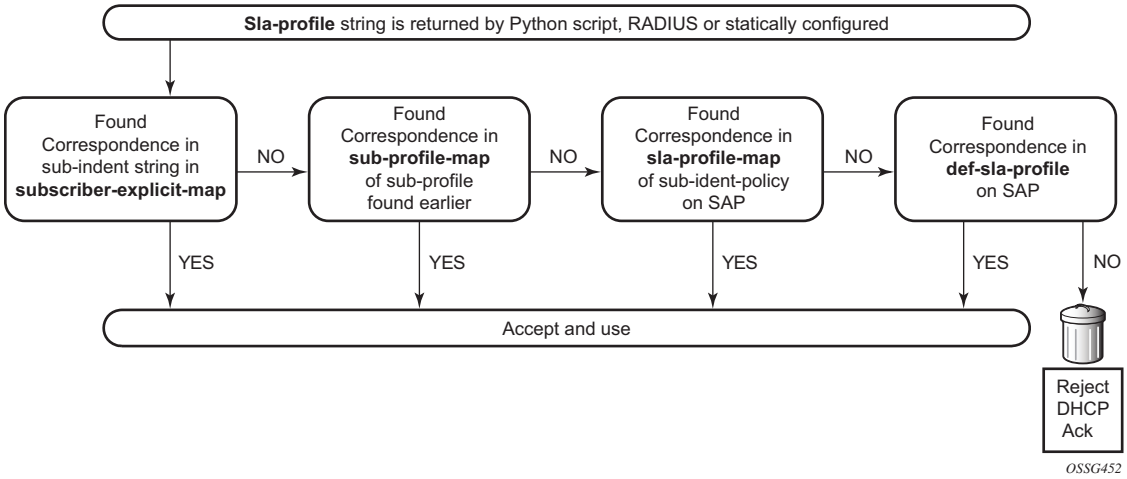

Note: Static configuration has priority over RADIUS configuration and RADIUS has priority over DHCP/Python scripts.

Note: Each host can have different SLA-profile, while sub-profile applies to whole subscriber. The last definition of sub-profile will force all previously defined hosts to change their sub-profile.

Bridged CO supports typical access node connection models, such as:

- One VLAN per service (ESM for subscriber differentiation and SAP for service)
- One VLAN per subscriber (SAP for subscriber differentiation and QoS flag for service)
- One VLAN per access node (ESM for subscriber differentiation and QoS flag for service)

Each of these modes has its pros and cons, but this is out of scope of this document.

This configuration guide focuses on configuration of one subscriber with three different hosts. VLAN per service is used as mode of subscriber aggregation and mixed RADIUS and DHCP as subscriber identification method. IP termination is done in IES service of BSR.

Correlation of BSA/BSR services and subscriber hosts is presented in Table 10.

Table 10 **Correlation of Hosts and BSA/BSR Services**

	BSA (Service/Features)	**BSR (Service/Features)**
Host-1 ca:00:0c:54:00:08	VPLS-100 • DHCP proxy server • SAP/SDP DHCP snoop • Sub-Ident origin via RADIUS • Sla/Sub-profiles via RADIUS • IP options via RADIUS	IES-100
Host-2 ca:01:08:10:00:08	VPLS-200 • SAP/SDP DHCP snoop • Sub-Ident origin through RADIUS • Sla/Sub-profiles through RADIUS • IP options through DHCP	IES-200 • DHCP relay
Host-3 ca:02:02:d0:00:08	VPLS-300 • SAP/SDP DHCP snoop • Sub-Ident origin through DHCP • Sla/Sub-profiles through DHCP • IP options through DHCP	IES-300 • DHCP relay

Different methods of authentication and address allocation were chosen for demonstration purposes. The customer is not limited to one method and can use a combination of methods as presented in this guide.

The following entities should be configured in advanced. Refer to the appropriate platform user guide for specific information. See Preface for a list of documents.

- Basic router configuration (interfaces, routing protocols, MPLS)

- External RADIUS server
- External/Local DHCP server

Configuration

A sample topology is presented in Figure 83.

Figure 83 **Sample Topology**

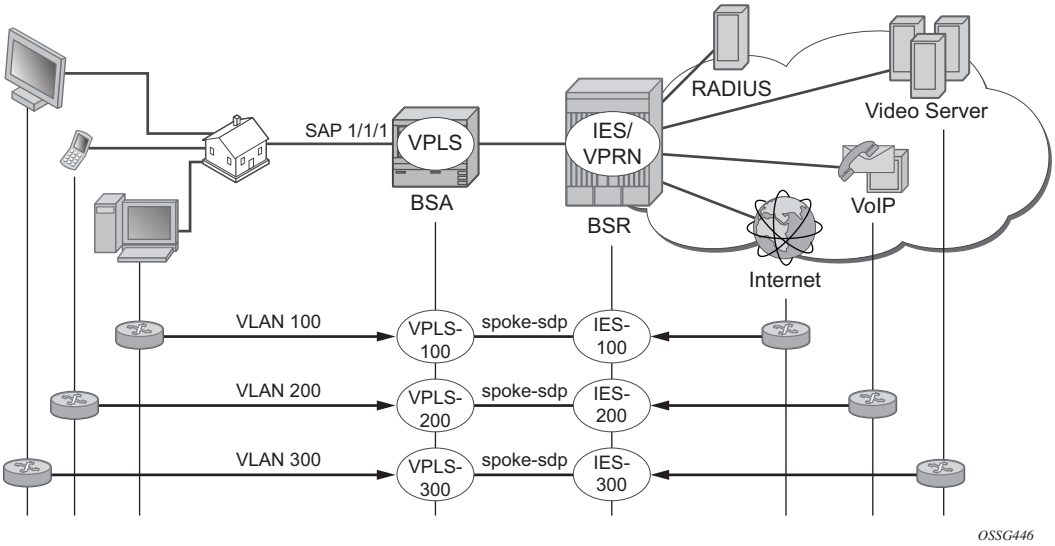

Bridged CO model requires certain techniques and features to be used on different nodes. Major methods are presented in Figure 84.

Figure 84 **Functionality of Each Node**

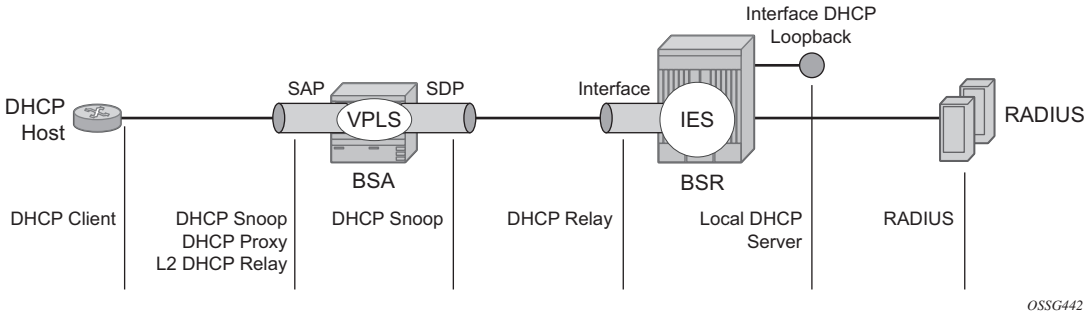

The following configuration steps are required:

Step 1. On BSA

1. Configure subscriber management profiles
    1. Configure sla profiles
    2. Configure subscriber profiles
    3. Configure subscriber identification policies
    4. Configure authentication and accounting policies if required
2. Configure VPLS service
    1. Configure split horizon group
    2. Configure SAP
        2.1 Configure anti-spoofing filters
        2.2 Configure DHCP snooping
        2.3 Configure optional parameters (lease split, L2 DHCP relay agent, etc.)
        2.4 In case of RADIUS authentication apply authentication policy
        2.5 Configure ESM
    3. Configure SDP
        3.1 Configure DHCP snooping

Step 2. On BSR

1. Configure IES service
    1. Configure IP interface
        1.1 Configure DHCP relay agent if required

Basic ESM Configuration on BSA

Subscriber management is enabled on BSA in Bridged CO model. A relevant configuration is presented below. SLA and subscriber profiles show the default configurations. The authentication policy appeals to RADIUS server 192.0.2.5. The subscriber identification policy is configured to use DHCP Option 254 to transfer custom attributes (subscriber-id, sla-profile, sub-profile, etc.)

```
A:BSA>config>subscr-mgmt# info
        authentication-policy "auth-policy-1" create
            password password-1
            radius-authentication-server
                router "management"
                server 1 address 192.0.2.5 secret ALU
            exit
            include-radius-attribute
                circuit-id
```

```
        remote-id
        nas-port-id
        nas-identifier
    exit
exit
sla-profile "sla-profile-1" create
exit
sla-profile "sla-profile-2" create
exit
sla-profile "sla-profile-3" create
exit
sla-profile "sla-profile-default" create
exit
sub-profile "sub-profile-1" create
exit
sub-profile "sub-profile-default" create
exit
sub-ident-policy "sub-ident-policy-1" create
    sub-profile-map
        use-direct-map-as-default
    exit
    sla-profile-map
        use-direct-map-as-default
    exit
    strings-from-option 254
exit
```

The **string-from-option 254** command is shared in-built dhcp-server of BSR. Using this option, the DHCP server could transmit subscriber identification options such the subscriber-id, sla-profile-string, and sub-profile-string.

BSA/BSR Configuration for Host-1 Operation

The test subscriber has three hosts. Host-1 gets all necessary information from RADIUS server.

Table 11 **BSA/BSR Configuration for Host-1 Operation**

	BSA (Service/Features)	**BSR (Service/Features)**
Host-1 ca:00:0c:54:00:08	VPLS-100 • DHCP proxy server • SAP/SDP DHCP snoop • Sub-Ident origin through RADIUS • Sla/Sub-profiles through RADIUS • IP options through RADIUS	IES-100

In this case BSA takes role of DHCP proxy with DHCP server emulation. DHCP snooping on the SAP must be enabled. Anti-spoofing filters on the SAP must be enabled. An authentication policy must be applied on the SAP.

```
        vpls 100 customer 1 create
            split-horizon-group "RSHG-1" residential-group create
            exit
--snip--
            sap 1/1/4:100 split-horizon-group "RSHG-1" create
                dhcp
                    snoop
                    lease-populate 400
                    proxy-server
                        emulated-server 10.0.1.253
                        no shutdown
                    exit
                    no shutdown
                exit
                authentication-policy "auth-policy-1"
                anti-spoof ip-mac
                sub-sla-mgmt
                    def-sub-id string "default-subscriber"
                    def-sub-profile "sub-profile-default"
                    def-sla-profile "sla-profile-default"
                    sub-ident-policy "sub-ident-policy-1"
                    no shutdown
                exit
            exit
            spoke-sdp 12:100 create
            exit
            no shutdown
        exit
```

On BSR IES-100, the service is configured with a pure IP interface, which plays role of DG for host-1.

```
ies 100 customer 1 create
    interface "int-host-1" create
        address 10.0.1.254/24
        spoke-sdp 21:100 create
        exit
    exit
    no shutdown
exit
```

BSA/BSR Configuration for Host-2 Operation

The test subscriber has three hosts. Host-2 gets subscriber-id and sla/sub-profiles information from RADIUS server and IP options from DHCP server.

Table 12 **BSA/BSR Configuration for Host-2 Operation**

	BSA (Service/Features)	**BSR (Service/Features)**
Host-2 ca:01:08:10:00:08	VPLS-200 • SAP/SDP DHCP snoop • Sub-Ident origin through RADIUS • Sla/Sub-profiles through RADIUS • IP options through DHCP	IES-200 • DHCP relay

DHCP snooping on the SAP and SDP must be enabled. Anti-spoofing filters on the SAP must be enabled.

```
        vpls 200 customer 1 create
            split-horizon-group "RSHG-1" residential-group create
            exit
--snip--
            sap 1/1/4:200 split-horizon-group "RSHG-1" create
                dhcp
                    snoop
                    lease-populate 400
                    no shutdown
                exit
                authentication-policy "auth-policy-1"
                anti-spoof ip-mac
                sub-sla-mgmt
```

```
            def-sub-id string "default-subscriber"
            def-sub-profile "sub-profile-default"
            def-sla-profile "sla-profile-default"
            sub-ident-policy "sub-ident-policy-1"
            no shutdown
        exit
    exit
    spoke-sdp 12:200 create
        dhcp
            snoop
        exit
    exit
    no shutdown
exit
```

On BSR IES-200, the service is configured with an IP interface which as the DG for Host-2. DHCP relay must be configured to transform broadcast DHCP discover message into unicast and send it to DHCP server for processing.

```
ies 200 customer 1 create
    interface "int-host-2" create
        address 10.0.2.254/24
        dhcp
            server 192.0.2.4
            trusted
            no shutdown
        exit
        spoke-sdp 21:200 create
        exit
    exit
    no shutdown
exit
```

BSA/BSR Configuration for Host-3 Operation

The test subscriber has three hosts. Host-3 receives all necessary information from the DHCP server.

Table 13 **BSA/BSR Configuration for Host-3 Operation**

	BSA (Service/Features)	**BSR (Service/Features)**
Host-3 ca:02:02:d0:00:08	VPLS-300 * SAP/SDP DHCP snoop * Sub-Ident origin through DHCP * Sla/Sub-profiles through DHCP * IP options through DHCP	IES-300 * DHCP relay

DHCP snooping on the SAP and SDP must be enabled. Anti-spoofing filters on the SAP must be enabled.

```
        vpls 300 customer 1 create
            split-horizon-group "RSHG-1" residential-group create
            exit
---- snip ---
            sap 1/1/4:300 split-horizon-group "RSHG-1" create
                dhcp
                    snoop
                    lease-populate 400
                    no shutdown
                exit
                anti-spoof ip-mac
                sub-sla-mgmt
                    def-sub-id string "default-subscriber"
                    def-sub-profile "sub-profile-default"
                    def-sla-profile "sla-profile-default"
                    sub-ident-policy "sub-ident-policy-1"
                    no shutdown
                exit
            exit
            spoke-sdp 12:300 create
                dhcp
                    snoop
                exit
            exit
            no shutdown
        exit
```

On BSR IES-300, the service is configured with IP interface, which plays role of DG for host-3. DHCP relay must be configured to transform broadcast DHCP discover message into unicast and send it to DHCP server for processing.

```
        ies 300 customer 1 create
            interface "int-host-3" create
                address 10.0.3.254/24
                dhcp
                    server 192.0.2.4
                    trusted
                    no shutdown
                exit
                spoke-sdp 21:300 create
                exit
            exit
            no shutdown
        exit
```

RADIUS Configuration Bridged CO

The username in the RADIUS access request is configurable and could be one of the following formats:

- mac — MAC Source Address of the DHCP DISCOVER message
- circuit-id — Taken from option 82 in the received DHCP message. If no circuit-id can be found, the DHCP-msg is rejected.
- tuple — Concatenation of MAC source address and circuit-ID
- ascii-converted-circuit-id — Identical to circuit-id, but the user name will be sent to the RADIUS server as a string of hex digits
- ascii-converted-tuple — Identical to tuple, but the circuit-id part of the user name will be sent to the RADIUS server as a string of hex digits

Note: Refer to IPv4 DHCP Hosts for detailed information about how to use different options.

```
A:BSA>config>subscr-mgmt>auth-plcy# user-name-format
  - user-name-format <format> [append domain-name]
  - no user-name-format

 <format>             : mac|circuit-id|tuple|ascii-converted-circuit-id|
                        ascii-converted-tuple
```

For simplicity, MAC format is used in this guide.

There are two hosts configured in the users file on RADIUS server:

- a:00:0c:54:00:08 — The mac address of host-1 host [VPLS/IES 100]. For host-1 all necessary parameters are returned: subscriber-id, sla/sub-profiles, IP parameters and lease time.
- a:01:08:10:00:08 — The mac address of host-2 host [VPLS/IES 200]. For host-2 only subscriber-id, sla/sub-profiles are returned, while ip parameters and lease time are returned from DHCP server.

```
ca:00:0c:54:00:08 Auth-Type := Local, User-Password == "password-1"
                    Alc-Subsc-ID-Str = "sub-id-1",
                    Alc-Subsc-Prof-Str == "sub-profile-1",
                    Alc-SLA-Prof-Str == "sla-profile-1",
                    Framed-IP-Address = 10.0.1.1,
                    Framed-IP-Netmask = 255.255.255.0,
                    Alc-Default-Router = 10.0.1.254,
                    Session-Timeout = 6000

ca:01:08:10:00:08 Auth-Type := Local, User-Password == "password-1"
                    Alc-Subsc-ID-Str = "sub-id-1",
```

```
                    Alc-Subsc-Prof-Str == "sub-profile-1",
                    Alc-SLA-Prof-Str == "sla-profile-2"
```

Local DHCP Server Configuration Bridged CO

In the setup local DHCP server is used with reference to local user database.

```
A:BSR>config>router>dhcp# info
            local-dhcp-server "dhcp-server-1" create
                user-db "user-db-1"
                pool "pool-1" create
                    subnet 10.0.2.0/24 create
                    exit
                    subnet 10.0.3.0/24 create
                    exit
                exit
                no shutdown
            exit
```

Note: Subnets must be configured, even if all IP parameters are returned from local user DB. Without this option, DHCP server do not return IP parameters.

The local user database is configured on BSR. Identification is done via MAC address of a host, which is taken from DHCP-Discover message. There are several possibilities to identify DHCP host. **match-list** command is used for this purpose.

```
*A:BSR>config>subscr-mgmt>loc-user-db>dhcp#                 match-list
  - no match-list
  - match-list <dhcp-match-type-1> [<dhcpmatch-type-2>...(up to 4 max)]

 <dhcp-match-type>    : circuit-id|mac|option60|remote-id|sap-id|service-id|
                        string|system-id
```

There are two hosts configured:

- a:01:08:10:00:08 — mac address of host-2 [VPLS/IES 200]. DHCP returns ip address, subnet mask and default route.
- a:02:02:d0:00:08 — mac address of host-3 [VPLS/IES 300]. DHCP returns all necessary parameters: subscriber-id, sla/sub-profiles and all ip options.

```
A:BSR>config>subscr-mgmt# info
        local-user-db "user-db-1" create
            dhcp
                match-list mac
```

```
            host "host-2" create
                host-identification
                    mac ca:01:08:10:00:08
                exit
                address 10.0.2.1
                options
                    subnet-mask 255.255.255.0
                    default-router 10.0.2.254
                exit
                no shutdown
            exit
            host "host-3" create
                host-identification
                    mac ca:02:02:d0:00:08
                exit
                address 10.0.3.1
                identification-strings 254 create
                    subscriber-id "sub-id-1"
                    sla-profile-string "sla-profile-3"
                    sub-profile-string "sub-profile-1"
                exit
                options
                    subnet-mask 255.255.255.0
                    default-router 10.0.3.254
                exit
                no shutdown
            exit
        exit
        no shutdown
    exit
```

Setup Procedures and Debugging

Subscriber/Host Verification

The initialization of all active subscribers and hosts can be shown using the **how service active-subscribers** command. Different options can be used to filter the output of the command.

```
A:BSA# show service active-subscribers
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber sub-id-1 (sub-profile-1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:1/1/4:100 - sla:sla-profile-1
-------------------------------------------------------------------------------
IP Address      MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.0.1.1        ca:00:0c:54:00:08 N/A       DHCP
```

```
-------------------------------------------------------------------------------
(2) SLA Profile Instance sap:1/1/4:200 - sla:sla-profile-2
-------------------------------------------------------------------------------
IP Address      MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.0.2.1        ca:01:08:10:00:08 N/A       DHCP
-------------------------------------------------------------------------------
(3) SLA Profile Instance sap:1/1/4:300 - sla:sla-profile-3
-------------------------------------------------------------------------------
IP Address      MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.0.3.1        ca:02:02:d0:00:08 N/A       DHCP
-------------------------------------------------------------------------------
Number of active subscribers : 1
===============================================================================
A:BSA#
```

Hierarchy of subscriber hosts is represented in a convenient form using following command.

```
A:BSA# show service active-subscribers hierarchy
===============================================================================
Active Subscriber hierarchy
===============================================================================
-- sub-id-1 (sub-profile-1)
   |
   |-- sap:1/1/4:100 - sla:sla-profile-1
   |   |
   |   |-- 10.0.1.1 - ca:00:0c:54:00:08 - N/A (DHCP)
   |   |
   |
   |-- sap:1/1/4:200 - sla:sla-profile-2
   |   |
   |   |-- 10.0.2.1 - ca:01:08:10:00:08 - N/A (DHCP)
   |   |
   |
   |-- sap:1/1/4:300 - sla:sla-profile-3
   |   |
   |   |-- 10.0.3.1 - ca:02:02:d0:00:08 - N/A (DHCP)
   |   |
```

Host-1 Setup Debug

The Host-1 setup process is shown in Figure 85.

Figure 85 **Host-1 Setup Process**

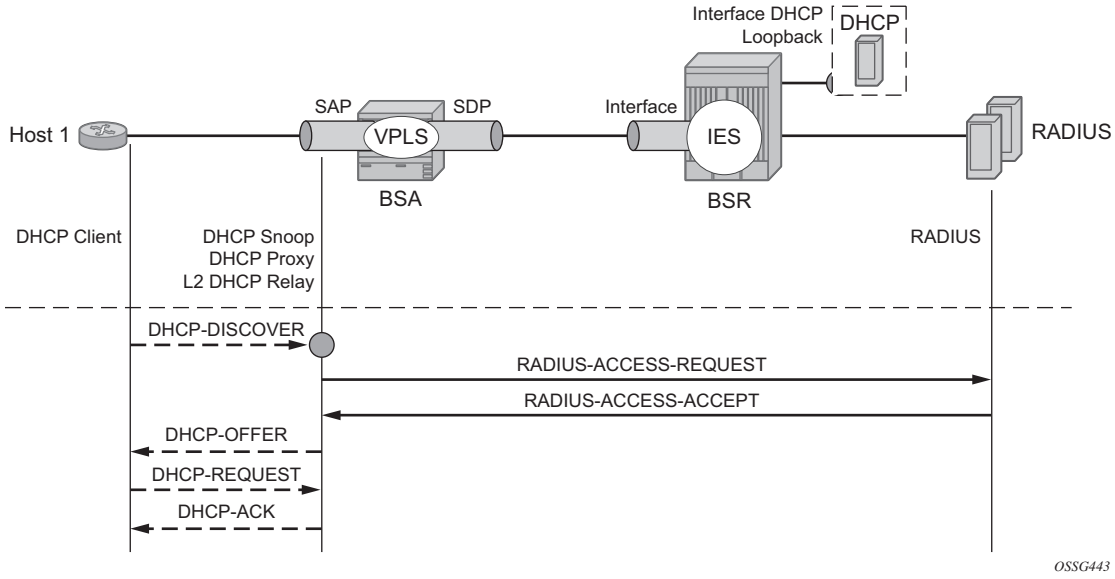

Host-1 sends DHCP discover message in VLAN 100 to BSA. BSA plays role of DHCP proxy server and transforms DHCP discover into RADIUS access-request message. After receiving RADIUS access-accept BSA transforms it to DHCP Ack message. Session setup process could be represented using debug commands:

```
A:BSA# debug service id 100 dhcp mode egr-ingr-and-dropped
A:BSA# debug service id 100 dhcp detail-level medium
A:BSA# debug radius detail
18 2009/12/15 06:31:56.63 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: RX DHCP Packet
   VPLS 100, SAP 1/1/4:100

   BootRequest to UDP port 67
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: ca:00:0c:54:00:08    xid: 0xd42

   DHCP options:
   [53] Message type: Discover
--snip--
"
19 2009/12/15 06:31:56.63 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Access-Request
  user ca:00:0c:54:00:08  policy auth-policy-1"
20 2009/12/15 06:31:56.63 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Transmit
  Access-Request(1) 192.0.2.5:1812 id 69  len 85
    USER NAME [1] 17 ca:00:0c:54:00:08
    PASSWORD [2] 16 lkhSVrFePQ0A0Xc4ZyMwMk
    NAS IP ADDRESS [4] 4 192.0.2.1
    NAS PORT TYPE [61] 4 Ethernet(15)
    NAS PORT ID [87] 9 1/1/4:100
```

```
    NAS IDENTIFIER [32] 3 BSA
"
21 2009/12/15 06:31:56.73 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Receive
  Access-Accept(2) id 69  len 108 from 138.203.18.79:1812
    VSA [26] 10 Alcatel(6527)
      SUBSC ID STR [11] 8 sub-id-1
    VSA [26] 15 Alcatel(6527)
      SUBSC PROF STR [12] 13 sub-profile-1
    VSA [26] 15 Alcatel(6527)
      SLA PROF STR [13] 13 sla-profile-1
    FRAMED IP ADDRESS [8] 4 10.0.1.1
    FRAMED IP NETMASK [9] 4 255.255.255.0
    VSA [26] 6 Alcatel(6527)
      DEFAULT ROUTER [18] 4 10.0.1.254
    SESSION TIMEOUT [27] 4 6000
"
22 2009/12/15 06:31:56.73 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: TX DHCP Packet
   VPLS 100, SAP 1/1/4:100

   BootReply to UDP port 68
   ciaddr: 0.0.0.0           yiaddr: 10.0.1.1
   siaddr: 10.0.1.253        giaddr: 0.0.0.0
   chaddr: ca:00:0c:54:00:08    xid: 0xd42

   DHCP options:
   [53] Message type: Offer
--snip--
"
23 2009/12/15 06:31:57.57 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: RX DHCP Packet
   VPLS 100, SAP 1/1/4:100

   BootRequest to UDP port 67
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: ca:00:0c:54:00:08    xid: 0xd42

   DHCP options:
   [53] Message type: Request
--snip--
"
24 2009/12/15 06:31:57.57 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: TX DHCP Packet
   VPLS 100, SAP 1/1/4:100

   BootReply to UDP port 68
   ciaddr: 0.0.0.0           yiaddr: 10.0.1.1
   siaddr: 10.0.1.253        giaddr: 0.0.0.0
   chaddr: ca:00:0c:54:00:08    xid: 0xd42

   DHCP options:
   [53] Message type: Ack
--snip
```

The number of snooped/forwarded/dropped/proxied DHCP packets can be checked using the **show service id 100 dhcp statistics** command.

```
A:BSA# show service id 100 dhcp statistics
===============================================
DHCP Statistics, service 100
===============================================
Client Packets Snooped               : 2
Client Packets Forwarded             : 0
Client Packets Dropped               : 0
Client Packets Proxied (RADIUS)      : 2
Client Packets Proxied (Lease-Split) : 0
Server Packets Snooped               : 0
Server Packets Forwarded             : 0
Server Packets Dropped               : 0
DHCP RELEASEs Spoofed                : 0
DHCP FORCERENEWs Spoofed             : 0
===============================================
A:BSA#
```

Connectivity of Host-1could be checked with the **show service id 100 subscriber-hosts** command. Different options can be used to filter output of the command.

```
A:BSA# show service id 100 subscriber-hosts detail
===============================================================================
Subscriber Host table
===============================================================================
Sap                    IP Address      MAC Address       PPPoE-SID Origin
  Subscriber
-------------------------------------------------------------------------------
1/1/4:100              10.0.1.1        ca:00:0c:54:00:08 N/A       DHCP
  sub-id-1
-------------------------------------------------------------------------------
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
-------------------------------------------------------------------------------
Number of subscriber hosts : 1
```

The DHCP lease state can be checked with the **show service id 100 dhcp lease-state** command. Different options can be used to filter output of a command.

```
A:BSA# show service id 100 dhcp lease-state detail
===============================================================================
DHCP lease states for service 100
===============================================================================
Service ID           : 100
IP Address           : 10.0.1.1
Client HW Address    : ca:00:0c:54:00:08
SAP                  : 1/1/4:100
Remaining Lifetime   : 01h33m41s
Persistence Key      : N/A
Sub-Ident            : "sub-id-1"
Sub-Profile-String   : "sub-profile-1"
```

```
SLA-Profile-String   : "sla-profile-1"
--snip--
Sub-Ident origin     : Radius
Strings origin       : Radius
Lease Info origin    : Radius
--snip--
Radius User-Name     : "ca:00:0c:54:00:08"
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:BSA#
```

Host-2 Setup Debug

Host-1 setup process is displayed in Figure 86.

Figure 86 **Host-2 Setup Process**

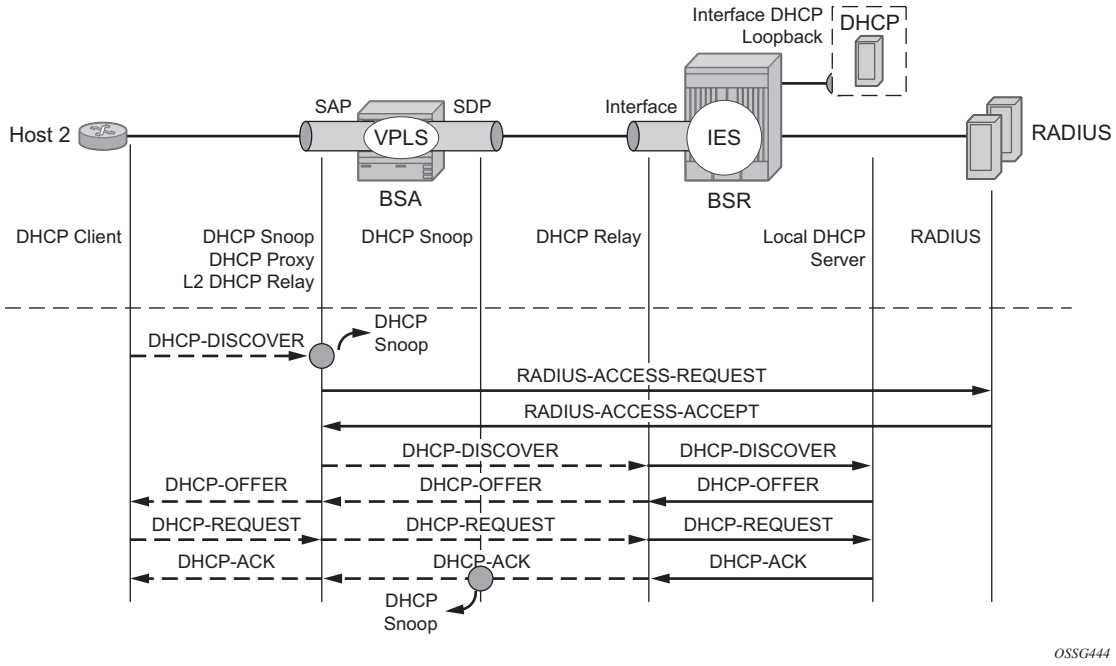

Host-2 sends DHCP discover message in VLAN 200 to BSA. Host-2 is authenticated through RADIUS and gets subscriber-id, sla/sub-profiles. DHCP Discover message is flooded in VPLS service and reaches IP interface on BSA, where DHCP relay is configured. Session setup process could be represented using debug commands:

```
A:BSA# debug service id 200 dhcp mode egr-ingr-and-dropped
A:BSA# debug service id 200 dhcp detail-level medium
A:BSA#
*A:BSA#
18 2009/12/15 13:00:36.28 UTC MINOR: DEBUG #2001 Base SVCMGR
```

```
"SVCMGR: RX DHCP Packet
   VPLS 200, SAP 1/1/4:200

   BootRequest to UDP port 67
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: ca:01:08:10:00:08    xid: 0xfc8

   DHCP options:
   [53] Message type: Discover
--snip--

19 2009/12/15 13:00:36.28 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Access-Request
  user ca:01:08:10:00:08  policy auth-policy-1"

20 2009/12/15 13:00:36.28 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Transmit
  Access-Request(1) 192.0.2.5:1812 id 80  len 85
    USER NAME [1] 17 ca:01:08:10:00:08
    PASSWORD [2] 16 .czdppt/0qAsqKqStbvnV.
    NAS IP ADDRESS [4] 4 192.0.2.1
    NAS PORT TYPE [61] 4 Ethernet(15)
    NAS PORT ID [87] 9 1/1/4:200
    NAS IDENTIFIER [32] 3 BSA
"

21 2009/12/15 13:00:36.34 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Receive
  Access-Accept(2) id 80  len 78 from 138.203.18.79:1812
    VSA [26] 10 Alcatel(6527)
      SUBSC ID STR [11] 8 sub-id-1
    VSA [26] 15 Alcatel(6527)
      SUBSC PROF STR [12] 13 sub-profile-1
    VSA [26] 15 Alcatel(6527)
      SLA PROF STR [13] 13 sla-profile-2
"
22 2009/12/15 13:00:36.34 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: TX DHCP Packet
   flooding in VPLS 200

   BootRequest to UDP port 67
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: ca:01:08:10:00:08    xid: 0xfc8

   DHCP options:
   [53] Message type: Discover
--snip--"

23 2009/12/15 13:00:36.35 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: RX DHCP Packet
   VPLS 200, spoke-sdp 12:200

   BootReply to UDP port 68
   ciaddr: 0.0.0.0           yiaddr: 10.0.2.1
   siaddr: 192.0.2.4         giaddr: 0.0.0.0
   chaddr: ca:01:08:10:00:08    xid: 0xfc8
```

```
   DHCP options:
   [53] Message type: Offer
--snip--"

24 2009/12/15 13:00:36.34 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: TX DHCP Packet
   VPLS 200, SAP 1/1/4:200

   BootReply to UDP port 68
   ciaddr: 0.0.0.0           yiaddr: 10.0.2.1
   siaddr: 192.0.2.4         giaddr: 0.0.0.0
   chaddr: ca:01:08:10:00:08    xid: 0xfc8

   DHCP options:
   [53] Message type: Offer
--snip--"

25 2009/12/15 13:00:36.46 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: RX DHCP Packet
   VPLS 200, SAP 1/1/4:200

   BootRequest to UDP port 67
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: ca:01:08:10:00:08    xid: 0xfc8

   DHCP options:
   [53] Message type: Request
--snip--"

26 2009/12/15 13:00:36.46 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: TX DHCP Packet
   flooding in VPLS 200

   BootRequest to UDP port 67
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: ca:01:08:10:00:08    xid: 0xfc8

   DHCP options:
   [53] Message type: Request
--snip--"

27 2009/12/15 13:00:36.47 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: RX DHCP Packet
   VPLS 200, spoke-sdp 12:200

   BootReply to UDP port 68
   ciaddr: 0.0.0.0           yiaddr: 10.0.2.1
   siaddr: 192.0.2.4         giaddr: 0.0.0.0
   chaddr: ca:01:08:10:00:08    xid: 0xfc8

   DHCP options:
   [53] Message type: Ack
--snip--"

28 2009/12/15 13:00:36.46 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: TX DHCP Packet
   VPLS 200, SAP 1/1/4:200
```

```
   BootReply to UDP port 68
   ciaddr: 0.0.0.0           yiaddr: 10.0.2.1
   siaddr: 192.0.2.4         giaddr: 0.0.0.0
   chaddr: ca:01:08:10:00:08    xid: 0xfc8

   DHCP options:
   [53] Message type: Ack
--snip--"

DHCP relay is enabled in service IES-200 on BSR.

A:BSR# debug router ip dhcp mode egr-ingr-and-dropped
A:BSR#
*A:BSR#
17 2009/12/15 13:00:36.34 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 6 (int-host-2),
   received DHCP Boot Request on Interface int-host-2 (1/1/2) Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: ca:01:08:10:00:08    xid: 0xfc8

   DHCP options:
   [53] Message type: Discover
--snip--"

18 2009/12/15 13:00:36.34 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
   transmitted DHCP Boot Request to 192.0.2.4 Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 10.0.2.254
   chaddr: ca:01:08:10:00:08    xid: 0xfc8

   DHCP options:
   [53] Message type: Discover
--snip--"

19 2009/12/15 13:00:36.35 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
   received DHCP Boot Reply on 192.0.2.4 Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.0.2.1
   siaddr: 192.0.2.4         giaddr: 10.0.2.254
   chaddr: ca:01:08:10:00:08    xid: 0xfc8

   DHCP options:
   [53] Message type: Offer
--snip--"

20 2009/12/15 13:00:36.35 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
```

```
instance 1 (Base), interface index 6 (int-host-2),
   transmitted DHCP Boot Reply to Interface int-host-2 (spoke-21:200) Port 68

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.0.2.1
   siaddr: 192.0.2.4         giaddr: 0.0.0.0
   chaddr: ca:01:08:10:00:08    xid: 0xfc8

   DHCP options:
   [53] Message type: Offer
--snip--"

21 2009/12/15 13:00:36.47 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 6 (int-host-2),
   received DHCP Boot Request on Interface int-host-2 (1/1/2) Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: ca:01:08:10:00:08    xid: 0xfc8

   DHCP options:
   [53] Message type: Request
--snip--"

22 2009/12/15 13:00:36.47 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
   transmitted DHCP Boot Request to 192.0.2.4 Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 10.0.2.254
   chaddr: ca:01:08:10:00:08    xid: 0xfc8

   DHCP options:
   [53] Message type: Request
--snip--"

23 2009/12/15 13:00:36.47 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
   received DHCP Boot Reply on 192.0.2.4 Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.0.2.1
   siaddr: 192.0.2.4         giaddr: 10.0.2.254
   chaddr: ca:01:08:10:00:08    xid: 0xfc8

   DHCP options:
   [53] Message type: Ack
--snip--"

24 2009/12/15 13:00:36.47 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 6 (int-host-2),
   transmitted DHCP Boot Reply to Interface int-host-2 (spoke-21:200) Port 68
```

```
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.0.2.1
   siaddr: 192.0.2.4         giaddr: 0.0.0.0
   chaddr: ca:01:08:10:00:08    xid: 0xfc8

   DHCP options:
   [53] Message type: Ack
--snip--
```

The number of snooped/forwarded/dropped/proxied DHCP packets could be checked using the **show service id 200 dhcp statistics** command.

```
A:BSA# show service id 200 dhcp statistics
=====================================================
DHCP Statistics, service 200
=====================================================
Client Packets Snooped               : 2
Client Packets Forwarded             : 2
Client Packets Dropped               : 0
Client Packets Proxied (RADIUS)      : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Snooped               : 2
Server Packets Forwarded             : 2
Server Packets Dropped               : 0
DHCP RELEASEs Spoofed                : 0
DHCP FORCERENEWs Spoofed             : 0
=====================================================
A:BSA#
```

The connectivity of Host-2 can be verified with the **show service id 200 subscriber-hosts** command. Different options can be used to filter output of the command.

```
A:BSA# show service id 200 subscriber-hosts detail
===============================================================================
Subscriber Host table
===============================================================================
Sap                    IP Address      MAC Address       PPPoE-SID Origin
  Subscriber
-------------------------------------------------------------------------------
1/1/4:200              10.0.2.1        ca:01:08:10:00:08 N/A       DHCP
  sub-id-1
-------------------------------------------------------------------------------
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-2
App Profile           : N/A
-------------------------------------------------------------------------------
Number of subscriber hosts : 1
===============================================================================
A:BSA#
```

DHCP lease state can be verified with the **show service id 200 dhcp lease-state** command. Different options can be used to filter output of the command.

```
A:BSA# show service id 200 dhcp lease-state detail
===============================================================================
DHCP lease states for service 200
===============================================================================
Service ID           : 200
IP Address           : 10.0.2.1
Client HW Address    : ca:01:08:10:00:08
SAP                  : 1/1/4:200
Remaining Lifetime   : 09d23h44m
Persistence Key      : N/A
Sub-Ident            : "sub-id-1"
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-2"
--snip--
Sub-Ident origin     : Radius
Strings origin       : Radius
Lease Info origin    : DHCP
--snip--
Radius User-Name     : "ca:01:08:10:00:08"
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:BSA#
```

Host-3 Setup Debug

The Host-3 setup process is presented in Figure 87.

Figure 87 **Host-3 Setup Process**

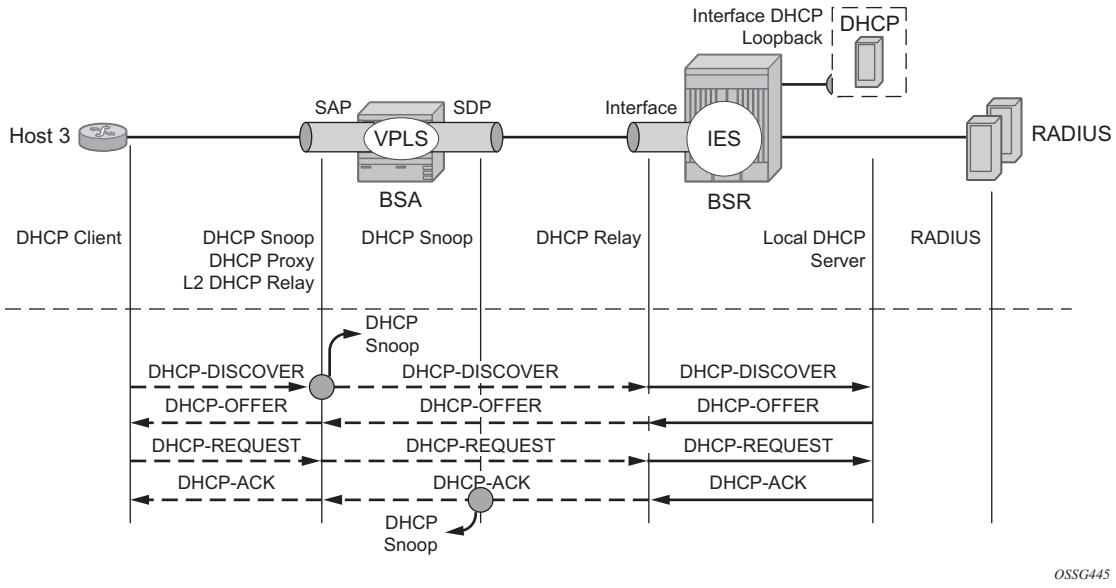

Host-3 sends DHCP a discover message in VLAN 300 to BSA. Host-3 receives all parameters from the DHCP server using pre-configured Option 254. A DHCP discover message is flooded in VPLS service and reaches IP interface on BSA where DHCP relay is configured. The session setup process can be represented using debug commands.

```
A:BSA# debug service id 300 dhcp mode egr-ingr-and-dropped
A:BSA# debug service id 300 dhcp detail-level medium
*A:BSA#
33 2009/12/15 13:02:34.39 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: RX DHCP Packet
   VPLS 300, SAP 1/1/4:300

   BootRequest to UDP port 67
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: ca:02:02:d0:00:08    xid: 0x2a6

   DHCP options:
   [53] Message type: Discover
--snip--
"
34 2009/12/15 13:02:34.38 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: TX DHCP Packet
   flooding in VPLS 300

   BootRequest to UDP port 67
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: ca:02:02:d0:00:08    xid: 0x2a6

   DHCP options:
   [53] Message type: Discover
--snip--
"
35 2009/12/15 13:02:34.40 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: RX DHCP Packet
   VPLS 300, spoke-sdp 12:300

   BootReply to UDP port 68
   ciaddr: 0.0.0.0           yiaddr: 10.0.3.1
   siaddr: 192.0.2.4         giaddr: 0.0.0.0
   chaddr: ca:02:02:d0:00:08    xid: 0x2a6

   DHCP options:
   [53] Message type: Offer
--snip--
"
36 2009/12/15 13:02:34.40 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: TX DHCP Packet
   VPLS 300, SAP 1/1/4:300

   BootReply to UDP port 68
   ciaddr: 0.0.0.0           yiaddr: 10.0.3.1
   siaddr: 192.0.2.4         giaddr: 0.0.0.0
   chaddr: ca:02:02:d0:00:08    xid: 0x2a6

   DHCP options:
```

```
   [53] Message type: Offer
--snip--
"
37 2009/12/15 13:02:34.52 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: RX DHCP Packet
   VPLS 300, SAP 1/1/4:300

   BootRequest to UDP port 67
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: ca:02:02:d0:00:08    xid: 0x2a6

   DHCP options:
   [53] Message type: Request
--snip--
"
38 2009/12/15 13:02:34.52 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: TX DHCP Packet
   flooding in VPLS 300

   BootRequest to UDP port 67
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: ca:02:02:d0:00:08    xid: 0x2a6

   DHCP options:
   [53] Message type: Request
--snip--
"
39 2009/12/15 13:02:34.53 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: RX DHCP Packet
   VPLS 300, spoke-sdp 12:300

   BootReply to UDP port 68
   ciaddr: 0.0.0.0           yiaddr: 10.0.3.1
   siaddr: 192.0.2.4         giaddr: 0.0.0.0
   chaddr: ca:02:02:d0:00:08    xid: 0x2a6

   DHCP options:
   [53] Message type: Ack
--snip--
"
40 2009/12/15 13:02:34.53 UTC MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: TX DHCP Packet
   VPLS 300, SAP 1/1/4:300

   BootReply to UDP port 68
   ciaddr: 0.0.0.0           yiaddr: 10.0.3.1
   siaddr: 192.0.2.4         giaddr: 0.0.0.0
   chaddr: ca:02:02:d0:00:08    xid: 0x2a6

   DHCP options:
   [53] Message type: Ack
--snip--
```

DHCP relay is enabled in service IES-300 on the BSR.

```
A:BSR# debug router ip dhcp mode egr-ingr-and-dropped
*A:BSR#
29 2009/12/15 13:02:34.39 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 7 (int-VoD),
   received DHCP Boot Request on Interface int-VoD (1/1/2) Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: ca:02:02:d0:00:08    xid: 0x2a6

   DHCP options:
   [53] Message type: Discover
--snip--
30 2009/12/15 13:02:34.39 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
   transmitted DHCP Boot Request to 192.0.2.4 Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 10.0.3.254
   chaddr: ca:02:02:d0:00:08    xid: 0x2a6

   DHCP options:
   [53] Message type: Discover
--snip--
"
31 2009/12/15 13:02:34.39 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
   received DHCP Boot Reply on 192.0.2.4 Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.0.3.1
   siaddr: 192.0.2.4         giaddr: 10.0.3.254
   chaddr: ca:02:02:d0:00:08    xid: 0x2a6

   DHCP options:
   [53] Message type: Offer
--snip--
"
32 2009/12/15 13:02:34.39 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 7 (int-VoD),
   transmitted DHCP Boot Reply to Interface int-VoD (spoke-21:300) Port 68

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.0.3.1
   siaddr: 192.0.2.4         giaddr: 0.0.0.0
   chaddr: ca:02:02:d0:00:08    xid: 0x2a6

   DHCP options:
   [53] Message type: Offer
--snip--
"
33 2009/12/15 13:02:34.53 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
```

```
instance 1 (Base), interface index 7 (int-VoD),
   received DHCP Boot Request on Interface int-VoD (1/1/2) Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: ca:02:02:d0:00:08    xid: 0x2a6

   DHCP options:
   [53] Message type: Request
--snip--
"
34 2009/12/15 13:02:34.53 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
   transmitted DHCP Boot Request to 192.0.2.4 Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 10.0.3.254
   chaddr: ca:02:02:d0:00:08    xid: 0x2a6

   DHCP options:
   [53] Message type: Request
--snip--
"
35 2009/12/15 13:02:34.53 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
   received DHCP Boot Reply on 192.0.2.4 Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.0.3.1
   siaddr: 192.0.2.4         giaddr: 10.0.3.254
   chaddr: ca:02:02:d0:00:08    xid: 0x2a6

   DHCP options:
   [53] Message type: Ack
--snip--
36 2009/12/15 13:02:34.53 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 7 (int-VoD),
   transmitted DHCP Boot Reply to Interface int-VoD (spoke-21:300) Port 68

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.0.3.1
   siaddr: 192.0.2.4         giaddr: 0.0.0.0
   chaddr: ca:02:02:d0:00:08    xid: 0x2a6

   DHCP options:
   [53] Message type: Ack
--snip--"
```

The number of snooped/forwarded/dropped/proxied DHCP packets can be verified with the using **show service id 300 dhcp statistics** command.

```
A:BSA# show service id 300 dhcp statistics
=====================================================
```

```
DHCP Statistics, service 300
==================================================
Client Packets Snooped               : 2
Client Packets Forwarded             : 2
Client Packets Dropped               : 0
Client Packets Proxied (RADIUS)      : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Snooped               : 2
Server Packets Forwarded             : 2
Server Packets Dropped               : 0
DHCP RELEASEs Spoofed                : 0
DHCP FORCERENEWs Spoofed             : 0
==================================================
A:BSA#
```

The connectivity of Host-3 can be verified with the **show service id 300 subscriber-hosts** command. Different options can be used to filter output of a command.

```
A:BSA# show service id 300 subscriber-hosts detail
===============================================================================
Subscriber Host table
===============================================================================
Sap                    IP Address      MAC Address       PPPoE-SID Origin
  Subscriber
-------------------------------------------------------------------------------
1/1/4:300              10.0.3.1        ca:02:02:d0:00:08 N/A       DHCP
  sub-id-1
-------------------------------------------------------------------------------
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-3
App Profile           : N/A
-------------------------------------------------------------------------------
Number of subscriber hosts : 1
===============================================================================
A:BSA#
```

The DHCP lease state could be checked with the **show service id 300 dhcp lease-state** command. Different options can be used to filter output of a command.

```
A:BSA# show service id 300 dhcp lease-state detail
===============================================================================
DHCP lease states for service 300
===============================================================================
Service ID           : 300
IP Address           : 10.0.3.1
Client HW Address    : ca:02:02:d0:00:08
SAP                  : 1/1/4:300
Remaining Lifetime   : 09d23h52m
Persistence Key      : N/A
Sub-Ident            : "sub-id-1"
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-3"
--snip--
Sub-Ident origin     : DHCP
Strings origin       : DHCP
Lease Info origin    : DHCP
--snip--
```

```
Radius User-Name     : ""
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:BSA#
```

Advanced Topics

Limiting Number of Subscribers

This topic is discussed in DHCP hosts. Refer to IPv4 DHCP Hosts for detailed information.

```
vpls 100 customer 1 create
    --snip--
    sap 1/1/4:100 split-horizon-group "RSHG-1" create
        --snip--
        sub-sla-mgmt
            --snip--
            multi-sub-sap 2
```

Limiting Number of Lease States

This topic is discussed in DHCP hosts. Refer to IPv4 DHCP Hosts for detailed information.

```
vpls 100 customer 1 create
    --snip--
    sap 1/1/4:100 split-horizon-group "RSHG-1" create
        dhcp
           lease-populate 400
```

Limiting Number of Host Per SLA-Profile

This topic is discussed in DHCP hosts. Refer to IPv4 DHCP Hosts for detailed information.

```
subscriber-mgmt
    sla-profile "sla-profile-1" create
        host-limit 1 [remove-oldest]
```

Subscriber Host Connectivity Verification

This topic is discussed in DHCP hosts. Refer to IPv4 DHCP Hosts for detailed information.

```
        vpls 100 customer 1 create
            sap 1/1/4:100 split-horizon-group "RSHG-1" create
                host-connectivity-verify source-ip 10.1.0.253 source-
mac 1e:54:ff:00:00:00 interval 1 action remove


A:BSA# show service id 100 host-connectivity-verify statistics
===============================================================================
Host connectivity check statistics
===============================================================================
Svc    SapId/          DestIp          Timestamp              Time since  Oper
Id     SdpId           Address         last-reply/conn-lost   Reply/Lost  State
-------------------------------------------------------------------------------
100    1/1/4:100       10.0.1.1        12/15/2009 09:04:06    0d 00:00:11 Up
-------------------------------------------------------------------------------
1 host-connectivity states : 1 Up / 0 Down / 0 Retry pending
===============================================================================
A:BSA#
```

Lease Split

This topic is discussed in DHCP hosts. Refer to IPv4 DHCP Hosts for detailed information.

```
        vpls 100 customer 1 create
            --snip--
            sap 1/1/4:100 split-horizon-group "RSHG-1" create
                dhcp
                    proxy-server
                       lease-time hrs 1
```

DHCP Option 82

This topic is discussed in DHCP hosts. Refer to IPv4 DHCP Hosts for detailed information.

Conclusion

This note provides configuration and troubleshooting commands for Bridged CO model.

DHCPv4 Server Basics

This chapter describes DHCPv4 server basics.

Topics in this chapter include:

Applicability

This chapter is applicable to SR OS routers and is based on SR OS release 14.0.R4.

Overview

The Dynamic Host Configuration Protocol for IPv4 (DHCPv4) provides a method for assigning addresses to hosts, and conveys additional configuration data to these hosts.

DHCPv4 allows for a flexible mapping of addresses to devices; for example, identified through their MAC address. While the DHCPv4 server owns and manages addresses organized in one or more pools, a DHCPv4 client obtains an address from the DHCPv4 server, which creates a lease for that client. This provides the client the right to use the address, and the server ensures that the address will not be assigned to other clients.

The DHCPv4 server implemented in SR OS has the following features:

- Address management. The DHCPv4 server keeps track of the used and unused addresses. For the used addresses, lease durations are maintained.
- Configuration parameter management. The DHCPv4 server stores parameters that are to be used by clients when they connect.
- Persistency. When enabled, the DHCPv4 server stores the leases on non-volatile storage so that the leases remain across potential node reboots.
- Failover capability. In dual-homed DHCPv4 server scenarios, a primary DHCPv4 server can take over the responsibility of a failing peer.

The DHCPv4 server failover capability is beyond the scope of this chapter.

In this chapter, when DHCP is mentioned, it implies DHCPv4.

Characteristics

IPv4 addresses and parameters are provided by the DHCP server through the Discover – Offer – Request – Acknowledge (DORA) message sequence as explained in the IPv4 DHCP Hosts chapter.

IPoE clients use DORA messages to communicate with the DHCP server via a relay agent. PPP clients use Link Control Protocol (LCP) and Internet Protocol Control Protocol (IPCP) to communicate with the router, and a local (DHCP) client manages the communication with the DHCP server. IPoE clients and PPP clients are also referred to as native clients and simulated clients, respectively.

When the DHCP server and the relay agent are physically located in the same SR OS node, the server is referred to as a local DHCP server; when they are in different nodes, the server is considered remote. Clients can obtain an address from a local, remote, or (external) third-party DHCP server.

A DHCP server can be used for IPoE users and PPP users simultaneously. A DHCP server must be hosted by a VPRN service or the base router. It can be accessed in either of the following ways; see Figure 88:

- When a DHCP user connects, the DORA message sequence running between the DHCP client and the DHCP server also passes through a relay agent, adding and removing user-defined options along the way. The relay agent and the DHCP server can be located in the same or in a different (remote) node.
- When a PPP user connects through LCP and IPCP on a service with an internal DHCP client (local client) configured, the local client manages the DORA communication toward the DHCP server, if the relay agent also has relaying enabled for PPP applications. The local client and the DHCP server can be located in the same or in a different (remote) node.
- When a PPP user connects through LCP and IPCP on a service with local address assignment configured, the DHCP server is accessed through an API-call. See the ESM SLAAC Prefix Assignment via Local Address Server chapter for an explanation of the local address assignment concept, which also applies to PPP.

Figure 88 **Accessing a DHCP server**

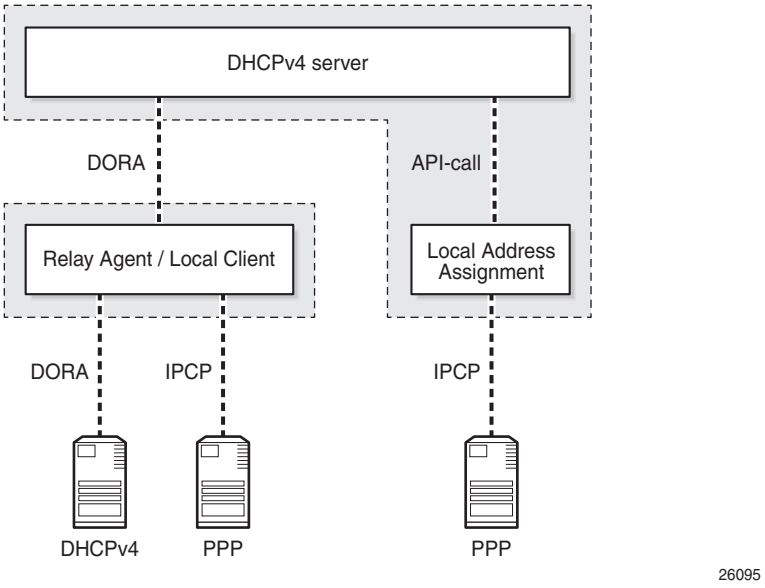

A DHCP server is supported for the routed CO model as well as for the bridged CO model.

A DHCP server can be hosted by the base router or a VPRN service, for public or private use. Because multiple VPRN services can coexist in a single node, each having its own DHCP server, overlapping address ranges are supported.

DHCP Lease

The DHCP server maintains the following data for every allocation request in a lease:

1. client-type (PPP or DHCP)
2. IP address
3. MAC address
4. lease state
5. option 82, if relevant
6. option 60 (vendor class identifier), if relevant
7. lease timer related data
8. persistence key, if applicable

A lease for a single client is in one of the following states:

1. offered: The IP address was offered to the client. The client still has to acknowledge the offer by sending a DHCP request.
2. stable: The IP address is now in use by the client.
3. force-renew-pending: The IP address is in use by the client, but the server sends a DHCP force-renew message to the client, because an option has changed at pool, subnet, or client (via LUDB) level.
4. remove-pending: The IP address is in use by the client, but the corresponding subnet range is deleted. The server sends a force-renew message to the client to force the client to reinitialize in order to get a new IP address.
5. held: The IP address has been used by the client but the lease has expired. The lease is now in the hold list so that the client can get the same IP address upon the next request for a lease.
6. internal: The IP address has been leased via local address assignment and is in use.
7. internal-orphan: The IP address has been leased via local address assignment and is in use. However, there is no configured subnet to which this lease belongs, because it has been removed or because this lease was installed through dual-homing.
8. internal-offered: The IP address has been offered via local address assignment, but the client has not acknowledged the offer yet.
9. internal-held: The IP address has been offered via local address assignment, but the lease is currently not active. The address is now in the hold list so that the same IP address can be offered to the same client upon request of a lease.
10. sticky: The IP address is reserved for the client and will remain reserved until the reservation for it is cleared. The client will get the same IP address upon the next request for a lease.

User Identification

The key to the leases managed by the DHCP server is configured at server level, and can be set to one of the following values (the default value being *mac-circuit-id*):

```
configure (router | service vprn <service-id>) dhcp local-dhcp-server <server-name>
   user-ident {client-id|circuit-id|mac|mac-circuit-id|remote-id}
```

The client ID is DHCP option 61; the circuit ID and the remote ID are sub-options 1 and 2 of DHCP option 82, respectively.

Setting *user-ident* to, for example, *circuit-id* can provide a CPE the same IP address regardless of its MAC address; thereby facilitating CPE replacement scenarios.

Lease Hold Time

The usual way for a DHCP client to indicate to the DHCP server it does not need its lease anymore is by sending a release message to the server; this is referred to as a solicited release.

However, when a client gets disconnected, or loses power, no release message is received by the server and the lease ultimately expires; this is referred to as an unsolicited release.

Without a lease hold timer, a lease is immediately deleted when the client sends the release message, or when the lease expires. The corresponding address is returned back to the pool of free addresses, and can be assigned to different clients. There is no guarantee that a client gets the same address again.

With a lease hold timer defined, a lease is not immediately deleted when the hold timer expires. Instead, the lease is put in the *held* or *internal-held* state. The lease is deleted when the hold timer expires, and the address is returned back to the pool. When the client connects, renews, or rebinds its lease before the hold timer expires, the client gets its previous lease again. There is no guarantee that the client gets the same address.

A lease hold timer can optionally be defined at the DHCP server level using the following command:

```
configure (router | service vprn <service-id>) dhcp local-dhcp-server <server-name>
    lease-hold-time [days <days>] [hrs <hours>] [min <minutes>] [sec <seconds>]

 <days>               : [0..7305]
 <hours>              : [0..23]
 <minutes>            : [0..59]
 <seconds>            : [0..59]
```

If delayed deletion is also required on reception of a release message (solicited release), use the following command:

```
configure (router | service vprn <service-id>) dhcp local-dhcp-server <server-name>
    lease-hold-time-for solicited-release
```

The same behavior can be applied to IPSec, but that is beyond the scope of this chapter.

Address Allocation for Sticky Leases

Sticky leases provide a static mapping between a hardware address and an IP address. This means that a particular device always gets the same IP address.

A sticky lease requires a host name. Identification of the host can be through a MAC address, a circuit ID, a remote ID, or a combination of these. The IP address must be in the range of the parenting pool.

```
tools perform router <router-id> dhcp local-dhcp-server <server-name>
    pool <pool-name> create-sticky-lease <hostname>
                                    [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>]
 <hostname>           : [32 chars max]
 <ieee-address>       : xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx
 <circuit-id>         : [253 chars max]
 <client-id>          : [255 chars max]
 <ip-address>         : a.b.c.d
 <circuit-id-hex-st*> : [0x0..0xFFFFFFFF...(max 506 hex nibbles)]
 <client-id-hex-str*> : [0x0..0xFFFFFFFF...(max 510 hex nibbles)]
```

As an alternative to sticky leases, an LUDB can also be used to provide a static mapping between a hardware address and an IP address. See the Local User Database for DHCPv4 Server chapter. However, hosts added to a local user database can only survive a reboot by saving the configuration through the **admin save** command.

Pool and Subnet Management

The DHCPv4 servers manage IPv4 addresses, subnets, and pools. These are hierarchically related to one another; see Figure 89.

Figure 89 **Addresses, Subnets, and Pools in a DHCPv4 Server**

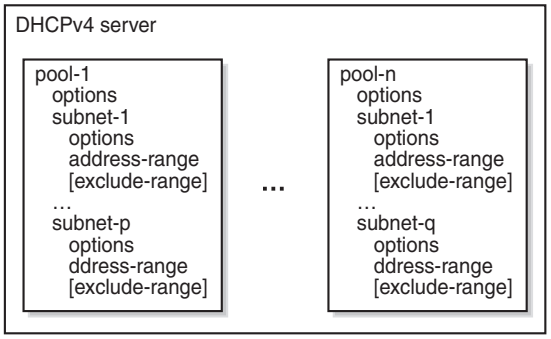

A subnet is identified by an IP address and a netmask, and defines:

- One or more address ranges – The ranges in the subnet that the server can allocate addresses from. Multiple address ranges cannot overlap.
- One or more exclude address ranges (optional) – A sub-range of the preceding range that the server will not allocate addresses from.
- Minimum-free – A notification is generated when the amount of free leases reaches this value (trap and log 99).
- Maximum-declined – Security counter measure, to prevent rogue clients from depleting the subnet. When this maximum value is reached, the oldest declined address will be returned to the pool.
- DHCP options:
  - default-router – up to four addresses can be defined
  - subnet-mask – subnet mask to be used by the clients
  - custom-options – additional options, when required

A pool is identified by name (maximum 32 characters), and defines:

- One or more subnets
- Min-lease-time – requests for a shorter lease time are set to this value; default is 10 min
- Max-lease-time – requests for a longer lease time are set to this value; default is 10 days
- Offer-time – a timer indicating how long an offer remains valid before the address offered is returned to the pool when no Request message is received (default 1 min)

- Minimum-free – a notification is generated when the amount of free leases reaches this value (trap and log 99), with an optional trap when all leases are used
- DHCP options:
  - dns-server – up to four DNS servers can be specified
  - domain-name – the domain to use for DNS resolution when clients provide unqualified host names.
  - lease-renew-time – defines when the client transitions to the renew state (T1)
  - lease-rebind-time – defines when the client transitions to the rebinding state (T2)
  - lease-time – the duration of time that the DHCP server grants to the DHCP client
  - netbios-name-server – defines up to four NetBIOS name servers
  - netbios-node-type – defines the NetBIOS node type (B, P, M, or H)
  - custom-option – additional options, when required

The options added by the DHCP server in response to an allocation request is a combination of the options provided by an LUDB (if applicable), subnet options, and pool options, in this sequence.

Lease Time

A DHCP client can request a specific lease time. The DHCP server checks for this value to be within the bounds as defined at pool level. If the requested lease time is out of bounds, it is set to either the minimum or the maximum value.

If a DHCP client does not request a specific lease time, the DHCP server takes the value from a matching LUDB entry, if available, or from the lease-time parameter defined at pool level, in this sequence. If the pool level lease time is not defined, the maximum lease time is used.

The best practice is to apply the following rule when defining values for the various timers:

```
lease-time > lease-rebind-time > lease-renew-time
```

However, the server does not check consistency of these timers, because the final values offered to the DHCP clients can come from various sources, which are out of the control of the DHCP server.

The local DHCP client always requests a lease time of 1 h to the server for PPP users connecting via the local client.

Address Allocation

When a request arrives at the DHCPv4 server, the server accesses the lease state database using the user ID as a key, checking for an existing lease. If a lease is already available, that lease is used.

Assuming that no lease is present in the lease state database yet, and that the server has a local user database attached, a matching entry is searched for in that local user database; see the Local User Database for DHCPv4 Server chapter.

In terms of address assignment, an LUDB attached to a DHCP server can return:

- an IP address – This (fixed) address is offered to the requester, where this address must not overlap with the address ranges configured in the local DHCP server.
- a Gi address – This address overrules any Gi address received from the requester.
- a pool name – A free address in one of the subnets in that pool is offered. Optionally, a secondary pool can be defined, which is used in case the primary pool is exhausted.
- use-gi-address [scope <subnet|pool>] – When the scope is set to subnet, the server offers an address from the subnet that includes the Gi address. When the scope is set to pool, the server offers an address from the subnet that includes the Gi address, or from the other subnets belonging to the same pool.
- use-pool-from-client [delimiter <delimiter>] – The pool name specified in the DHCP client message options and added by the relay agent is used. A free address in one of the subnets in that pool is offered. If two pools are available, the configured delimiting character identifies the splitting-point to find the names of both pools.

If a unique address is found in the LUDB, that address is offered by the server to the requester.

For the general address allocation flow, see Figure 90. The Local User Database for DHCPv4 Server chapter applies when an LUDB is attached to the DHCP server.

Figure 90 **General Address Allocation for DHCP**

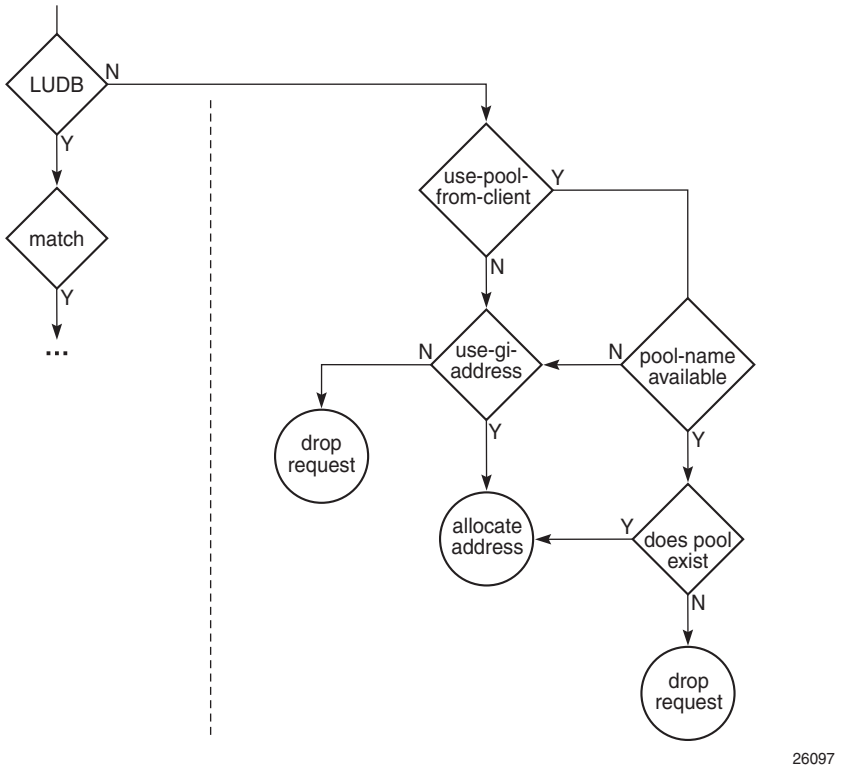

Two additional parameters are available at the server level, controlling which pool and subnet an address is taken from, as follows:

```
[no] use-pool-from-client [delimiter <delimiter>]
    <delimiter>          : [1 chars max]

[no] use-gi-address [scope <scope>]
    <scope>              : subnet|pool
```

With a requester-provided pool name and *use-pool-from-client* active, the server checks for that pool to exist before selecting a free address from one of the subnets in that pool.

With a requester-provided Gi address and *use-gi-address scope subnet* active, a free address is taken from the subnet that includes the Gi address. With *use-gi-address scope pool*, another subnet in the pool is used if the original subnet is exhausted.

The following rules apply to the DHCP server address allocation flow:

- Assume a DHCP server with an LUDB applied, and *use-gi-address* active:

  - A host lookup failure will not result in the request being dropped. The server sends an offer using an address selected based on the Gi address.
  - A successful host lookup, but returning a non-existent pool name, results in the server dropping the request, so no offer is sent.
- Assume a DHCP server without an LUDB applied, but with u*se-pool-from-client* and *use-gi-address* active:
  - A requester not providing a pool name results in the server sending an offer using an address selected based on the Gi address.
  - A requester providing a non-existent pool name results in the server dropping the request, so no offer is sent.

Therefore, *use-pool-from-client* takes precedence over *use-gi-address*. The DHCP server selects an address from a pool if that pool exists. If no pool name is provided to the DHCP server, address selection is based on the Gi address, when allowed through the *use-gi-address* directive.

The pool name provided by a relay agent can be a concatenation of two pool names, where the delimiter character is used to split the string apart in the original pool names.

Subnet Draining

When a subnet is put in the drained state through the drain command, no new leases can be assigned from that subnet. Existing leases are cleaned up upon renewal or rebinding of the client. This is useful in renumbering scenarios; see the Configuration section for an example.

Force Renew

Parameter force-renews enables DHCP servers to issue DHCP force-renew messages to stable clients, informing them about a configuration change.

With *force-renews* enabled, the server does not need to wait for a client to pass through its renew or rebind sequence to provide the reconfigured options, speeding up the configuration change.

Changes can be applied at the LUDB-level, subnet level, or pool-level.

DHCP Server Persistency

The DHCP protocol does not have a keep-alive mechanism to detect unavailability. Without precaution, a node reboot causes the loss of the DHCP lease state. Because DHCP clients only attempt a reinitialization sequence after expiration of the lease timer, service outages could become unacceptably long.

The DHCP server lease state can be made persistent across reboots. The lease state is then restored from a persistency file stored on the compact flash. Therefore, DHCP clients will only lose connectivity for the duration of the reboot, and no renew or rebind is needed.

Configuration

Starting a DHCP server up in an SR OS environment requires following steps:

- Configure the DHCP server.
- Configure the interfaces for the DHCP server to listen on.
- Configure one or more relay agents.

The baseline configuration used in this chapter is shown in Figure 91.

Figure 91 **Baseline Service Configuration**

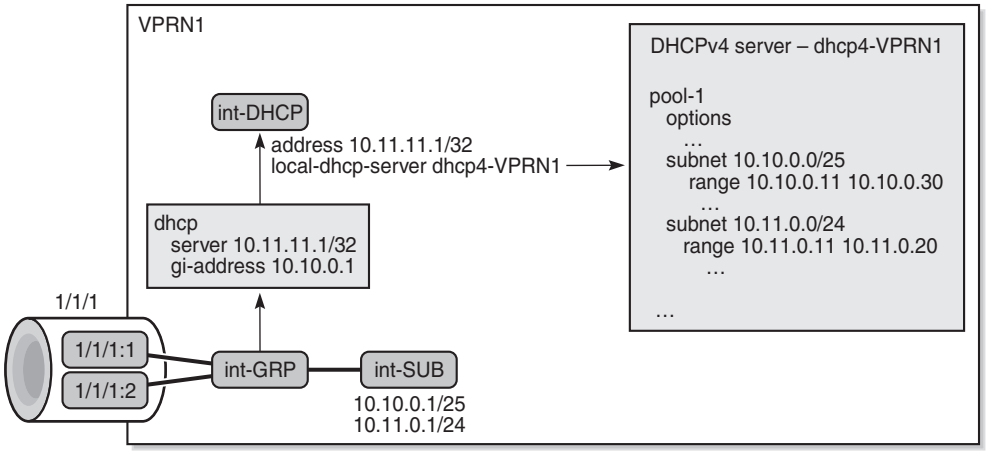

Configure the DHCP Server

One or more DHCP servers can be configured in the base router or in any routed service. VPRN 1 from Figure 91 has a single DHCP server named *dhcp4-VPRN1*, with two pools: *pool-1* and *pool-2*. The first pool contains two subnets; the second pool contains a single subnet.

Address selection is primarily based on a pool name (*use-pool-from-client*), and secondarily on the Gi address with the scope set to pool (*use-gi-address scope pool*). This means that address selection will be Gi address-based, if no pool name is provided to the server. Having the scope set to *pool* enables the server to allocate addresses from other subnets within the same pool.

Different options and custom options are defined at different levels. All subnets include an address range. Subnet 10.10.0.0/25 also has an exclusion range, as follows:

```
configure
    service
        vprn 1 customer 1 create
            dhcp
                local-dhcp-server dhcp4-VPRN1 create
                    use-gi-address scope pool
                    use-pool-from-client
                    pool "pool-1" create
                        options
                            dns-server 1.1.1.1 1.1.2.2
                            lease-time hrs 2
                            custom-option 150 address 1.1.1.1
                        exit
                        subnet 10.10.0.0/25 create
                            options
                                subnet-mask 255.255.255.128
                                default-router 10.10.0.1 10.10.0.2
                                custom-option 130 string "MyOption1"
                            exit
                            exclude-addresses 10.10.0.61 10.10.0.70
                            address-range 10.10.0.11 10.10.0.126
                        exit
                        subnet 10.11.0.0/24 create
                            options
                                subnet-mask 255.255.255.0
                                default-router 10.11.0.1
                                custom-option 130 string "MyOption2"
                            exit
                            address-range 10.11.0.11 10.11.0.20
                        exit
                    exit
                    pool "pool-2" create
                        subnet 10.20.0.0/16 create
                            options
                                subnet-mask 255.255.0.0
                                default-router 10.20.0.1
                            exit
```

```
                            address-range 10.20.0.21 10.20.0.120
                        exit
                    exit
                    no shutdown
                exit
            exit
        exit
    exit
exit
```

Configure the DHCP Interface

The DHCP server needs to be listening on one or more interfaces. In the example from Figure 90, the DHCP server is associated with interface *int-DHCP*, with loopback address 10.11.11.1, as follows. The DHCP server cannot be applied to a group interface.

```
configure
    service
        vprn 1 create
            interface "int-DHCP" create
                address 10.11.11.1/24
                local-dhcp-server "dhcp4-VPRN1"
                loopback
            exit
        exit
    exit
exit
```

Configure Relay Agents

The configuration of the DHCP server must align with the configuration of the relay agents for the server to assign addresses correctly. For example, defining the server to allocate addresses based on a pool name, but not providing a pool name toward the server, might not provide the expected result, because this will not necessarily lead to addresses being assigned and offered to clients.

The DHCP relay agent is configured in the dhcp context, as follows:

- *gi-address* – the gateway IPv4 address used by the relay agent
- *server* – up to 8 DHCP servers can be defined by their IPv4 address; only 10.11.11.11 is used in this example
- *client-applications dhcp ppp* – the DHCP server will allocate addresses for DHCP and PPP clients

- *option* – the options added/removed to/from messages toward the server. In the example, the circuit-id, the remote-id, and the pool-name are added.
- *trusted* – this parameter ensures that DHCP messages with option 82 included and the gi-address set to zero are being processed instead of being dropped

```
configure
    service
        vprn 1 customer 1 create
            subscriber-interface "int-SUB" create
                group-interface "int-GRP" create
                    dhcp
                        option
                            action replace
                            circuit-id
                            remote-id
                            vendor-specific-option
                                pool-name
                            exit
                        exit
                        server 10.11.11.1
                        lease-populate 100
                        client-applications dhcp ppp
                        gi-address 10.10.0.1
                        no shutdown
                    exit
                exit
            exit
        exit
    exit
exit
```

Configure DHCP Server Persistency

The following configuration stores the DHCP server lease-state persistency file on cf1:

```
configure
    system
        persistence
            dhcp-server
                location cf1:
            exit
        exit
    exit
exit
```

The persistency file is pre-allocated, providing space for the maximum number of allowed leases, which avoids file system space issues during normal operation, as follows:

```
*A:PE1>file cf1:\ # dir
```

```
Volume in drive cf1 on slot A has no label.

Volume in drive cf1 on slot A is formatted as FAT32

Directory of cf1:\

09/19/2016  04:29p      <DIR>          .ssh/
09/21/2016  01:58p           248513024 dhcp_serv.006
09/21/2016  01:58p             5825024 dhcp_serv.i06
               2 File(s)              254338048 bytes.
               1 Dir(s)              7759888384 bytes free.

*A:PE1>file cf1:\ #
```

A message is issued to log-id 99 to indicate that the persistence file is ready for use, as follows:

```
*A:PE1# show log log-id 99

===============================================================================
Event Log 99
===============================================================================
Description : Default System Log
Memory Log contents  [size=500   next event=10722  (wrapped)]
10721 2016/09/21 12:44:58.24 CEST WARNING: SYSTEM #2037 Base dhcp-
server Persistence Report
"Persistency event: dhcp-server persistence file ready for use"
```

The **tools dump persistence summary** command provides persistency information. The following example shows that the file cf1:\dhcp_serv.006 is used for storing persistency records for the DHCP server:

```
*A:PE1# tools dump persistence summary
===============================================================================
Persistence Summary on Slot A (active)
===============================================================================
Client           Location               #Registrations   File State
                 Avg Nr Fragments       #Entries         Subsystem State
                 File Fill Level        #Entries Queued
-------------------------------------------------------------------------------
dhcp-server      cf1:\dhcp_serv.006     5                ACTIVE
                 1.0                    5                OK
                 0%                     0
-------------------------------------------------------------------------------
Total for cf1:    3% in use
-------------------------------------------------------------------------------

*A:PE1#
```

Persistency records are identified using the persistence key. This key is part of the lease state. The following command shows the persistence key for lease 10.11.0.14:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" leases 10.11.0.14 detail

===============================================================================
```

```
Lease for DHCP server dhcp4-VPRN1 router 1
===============================================================================
IP-address                    : 10.11.0.14
Lease-state                   : stable
Lease started                 : 2016/10/17 15:28:36

--- snipped ---

User-db Address Type          : N/A
Persistence Key               : 0x00000004
Lease Remaining Hold Time     : 0h0m0s

===============================================================================
*A:PE1#
```

The DHCP server lease state records can be shown using the following command. This example shows the record for key 0x00000004:

```
*A:PE1# tools dump persistence dhcp-server record 0x00000004
----------------------------------------
Persistence Record
----------------------------------------
Client      : dhcp-server
Persist-Key : 0x00000004
Filename    : cf1:\dhcp_serv.006
Entries     : Index  FedHandle  Last Update               Action Valid
              004289 0x00000079 2016/10/17 14:48:46 (UTC) UPDATE Yes
Data        : 151 bytes

 type       : V4 lease
 service Id : 1
 server     : dhcp4-VPRN1
 IP         : 10.11.0.14
 MAC        : 00:00:00:01:01:03
 XID        : 0x00000020
 state      : stable
 lease mode : ET
 start time : 2016/10/17 13:28:36 (UTC)
 last renew : 2016/10/17 14:48:46 (UTC)
 expires    : 2016/10/17 14:58:46 (UTC)
 failctrl   : local
 opt60 len  : 0
 opt61 len  : 0
 opt82 len  : 0
 sticky name:
*A:PE1#
```

DHCP server lease state persistency is typically used together with subscriber management persistency if the DHCP server and subscriber management functions are managed by the same network node; see the IPv4 DHCP Hosts chapter.

Configure a Sticky Lease

The following command creates a sticky lease with name me-010101, using MAC address 00:00:00:01:02:02 and IP address 10.11.0.20:

```
*A:PE1# tools perform router 1 dhcp local-dhcp-server "dhcp4-VPRN1" pool "pool-
1" create-sticky-lease me-010202 mac 00:00:00:01:02:02 requested-ip-
address 10.11.0.20

===============================================================================
Sticky lease creation result
===============================================================================
Result                  : Success
IP-address              : 10.11.0.20
Lease-state             : sticky
Lease started           : 2016/10/17 17:07:00
Remaining LifeTime      : N/A
Sticky-lease Host Name  : me-010202
MAC address             : 00:00:00:01:02:02
Persistence Key         : N/A
===============================================================================
*A:PE1#
```

No user database may be assigned to the DHCP server to create sticky leases.

A **clear** command can be used to delete a sticky lease, as follows:

```
clear router 1 dhcp local-dhcp-server "dhcp4-VPRN1" sticky-leases hostname "me-
010202"
```

Operation and Verification

The following command shows all DHCP servers defined in the system. The maximum and active number of leases are shown. The router and services where the DHCP servers are hosted are listed, together with the server name and an indication whether this server is in- or out-of-service.

```
*A:PE1# show router dhcp servers all

===============================================================
Overview of DHCP Servers
===============================================================
Active Leases:      5
Maximum Leases:     159744

Router              Server                          Admin State
---------------------------------------------------------------
Service: 1          dhcp4-VPRN1                     inService
===============================================================
*A:PE1#
```

The following command shows all leases currently allocated by DHCP server dhcp4-VPRN1 in VPRN-1. In this example, the leases for the DHCP and PPP clients are all “stable”. Sticky leases are always shown, even when they are not online.

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" leases

===============================================================================
Leases for DHCP server dhcp4-VPRN1 router 1
===============================================================================
IP Address      Lease State       Mac Address       Remaining   Clnt  Fail
  PPP user name/Opt82 Circuit Id                    LifeTime    Type  Ctrl
  User-db/Sticky-lease Hostname
-------------------------------------------------------------------------------
10.10.0.11      stable            00:00:00:01:01:01 0h9m16s     dhcp  local

10.10.0.12      stable            00:00:00:01:01:02 0h7m36s     dhcp  local

10.11.0.14      stable            00:00:00:01:01:03 0h9m9s      dhcp  local

10.11.0.17      stable            00:00:00:00:00:33 0h59m55s    ppp   local
  PE1|1|int-GRP|1/1/1:1
10.11.0.20      sticky            00:00:00:01:02:02 N/A         dhcp  N/A

  me-010202
-------------------------------------------------------------------------------
5 leases found
===============================================================================
*A:PE1#
```

The following command shows the leases on the same server allocated from the 10.11.0.0/24 subnet:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" leases 10.11.0.0/24

===============================================================================
Leases for DHCP server dhcp4-VPRN1 router 1
===============================================================================
IP Address      Lease State       Mac Address       Remaining   Clnt  Fail
  PPP user name/Opt82 Circuit Id                    LifeTime    Type  Ctrl
  User-db/Sticky-lease Hostname
-------------------------------------------------------------------------------
10.11.0.14      stable            00:00:00:01:01:03 0h7m31s     dhcp  local

10.11.0.18      stable            00:00:00:00:00:33 0h59m40s    ppp   local
  PE1|1|int-GRP|1/1/1:1
10.11.0.20      sticky            00:00:00:01:02:02 N/A         dhcp  N/A

  me-010202
-------------------------------------------------------------------------------
3 leases found
===============================================================================
*A:PE1#
```

The following command shows the details of a single lease:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" leases 10.11.0.18/
```

```
32 detail

===============================================================================
Lease for DHCP server dhcp4-VPRN1 router 1
===============================================================================
IP-address                   : 10.11.0.18
Lease-state                  : stable
Lease started                : 2016/10/17 17:15:57
Last renew                   : N/A
Remaining LifeTime           : 0h57m55s
Remaining Potential Exp. Time: 0h0m0s
Sticky-lease Host Name       : N/A
MAC address                  : 00:00:00:00:00:33
Xid                          : 0x8bf01670
Failover Control             : local
Client Type                  : ppp
User-db Host Name            : N/A
User-db Address Type         : N/A
Persistence Key              : 0x00000005
Opt82 Hex Dump               : (length=71)
                             : 52 45 01 15 50 45 31 7c 31 7c 69 6e 74 2d 47 52
                             : 50 7c 31 2f 31 2f 31 3a 31 02 06 00 00 00 00 00
                             : 33 09 24 00 00 19 7f 1f 02 06 00 00 00 00 00 33
                             : 06 01 01 01 03 50 45 31 03 04 00 00 00 01 04 07
                             : 31 2f 31 2f 31 3a 31
Opt82 Circuit Id             : PE1|1|int-GRP|1/1/1:1
Opt82 Remote Id              : (hex) 00 00 00 00 00 33
Opt82 VS System              : PE1
Opt82 VS Clnt MAC            : 00:00:00:00:00:33
Opt82 VS Service             : (hex) 00 00 00 01
Opt82 VS SAP                 : 1/1/1:1
Opt82 VS String              :
Opt82 VS PPPoE Session ID    :
Opt60 Hex Dump               : (length=10)
                             : 41 4c 55 37 58 58 58 53 42 4d
Lease Remaining Hold Time    : 0h0m0s

===============================================================================
*A:PE1#
```

Troubleshooting

The following command shows summary data for the DHCP server:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" summary
===============================================================================
DHCP server dhcp4-VPRN1  router 1
===============================================================================
Admin State            : inService
Operational State      : inService
Persistency State      : ok
User Data Base         : N/A
Use gateway IP address : enabled (scope pool)
Use pool from client   : enabled
Send force-renewals    : disabled
```

```
Creation Origin        : manual
Lease Hold Time        : 0h10m0s
Lease Hold Time For    : (Not specified)
User-ident             : mac-circuit-id

Failover Admin State   : outOfService
Failover Oper State    : shutdown
Failover Persist Key   : 0x00000003
Administrative MCLT    : 0h10m0s
Operational MCLT       : 0h10m0s
Startup wait time      : 0h2m0s
Partner down delay     : 23h59m59s
  Ignore MCLT          : disabled

-------------------------------------------------------------------------------
Pool name : pool-1
-------------------------------------------------------------------------------
Failover Admin State   : outOfService
Failover Oper State    : shutdown
Failover Persist Key   : 0x00000001
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
-------------------------------------------------------------------------------
10.10.0.0/25           0        0%   2        0        0        0        N
10.11.0.0/24           241      98%  3        0        0        0        N
Totals for pool        241      97%  5        0        0        0
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Pool name : pool-2
-------------------------------------------------------------------------------
Failover Admin State   : outOfService
Failover Oper State    : shutdown
Failover Persist Key   : 0x00000007
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
-------------------------------------------------------------------------------
10.20.0.0/16           100      100% 0        0        0        0        N
Totals for pool        100      100% 0        0        0        0
-------------------------------------------------------------------------------

Totals for server      341      98%  5        0        0        0

-------------------------------------------------------------------------------
Interface associations
Interface                       Admin
-------------------------------------------------------------------------------
int-VPRN1-DHCPv4                Up
```

```
-------------------------------------------------------------------------------
Local Address Assignment associations
Group interface                  Admin
-------------------------------------------------------------------------------
===============================================================================
*A:PE1#
```

The following command shows DHCP server statistics:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" server-stats

===============================================================================
Statistics for DHCP Server dhcp4-VPRN1 router 1
===============================================================================
Rx Discover Packets           : 2449
Rx Request Packets            : 12752
Rx Release Packets            : 53
Rx Decline Packets            : 0
Rx Inform Packets             : 0

Tx Offer Packets              : 177
Tx Ack Packets                : 1184
Tx Nak Packets                : 63
Tx Forcerenew Packets         : 58

Client Ignored Offers         : 0
Leases Timed Out              : 2

Dropped Bad Packet            : 11205
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       : 5
Dropped Not Serving Pool      : 2297
Dropped Invalid User          : 0
Dropped Overload              : 0
Dropped Persistence Overload  : 0
Dropped Generic Error         : 0
Dropped Destined To Other     : 0
Dropped Address Unavailable   : 300
Dropped Max Leases Reached    : 0
Dropped Server Shutdown       : 0
Dropped No Subnet For Fixed IP: 0
Dropped Duplicate From Diff GI: 0
Dropped busy primary audit    : 0
Dropped transmission failed   : 0

Rx Internal Requests          : 0
Rx Internal Releases          : 0
Dropped Internal w/LUDB       : 0
Dropped Internal w/Failover   : 0
Dropped Internal w/Conflicts  : 0

Failover statistics
-------------------------------------------------------------------------------
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
===============================================================================
*A:PE1#
```

The following command shows extended server statistics:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" pool-ext-stats

===============================================================================
Extended pool statistics for server "dhcp4-VPRN1"
===============================================================================
                           Current         Peak            Peak Timestamp
-------------------------------------------------------------------------------
Pool                       pool-1
Local:
  Offered Leases           0               1               10/17/2016 17:15:57
  Stable Leases            5               5               10/17/2016 17:15:57
  Provisioned Addresses    246
  Used Addresses           5               5               10/17/2016 17:21:24
  Free Addresses           241             241             10/17/2016 17:21:24
  Used Pct                 3               3               10/17/2016 17:21:24
  Free Pct                 97              97              10/17/2016 17:21:24
Last Reset Time                                            10/17/2016 15:26:31
-------------------------------------------------------------------------------
Pool                       pool-2
Local:
  Offered Leases           0               0               10/17/2016 17:22:15
  Stable Leases            0               0               10/17/2016 17:22:15
  Provisioned Addresses    100
  Used Addresses           0               0               10/17/2016 17:22:15
  Free Addresses           100             100             10/17/2016 17:22:15
  Used Pct                 0               0               10/17/2016 17:22:15
  Free Pct                 100             100             10/17/2016 17:22:15
Last Reset Time                                            10/17/2016 17:22:15
-------------------------------------------------------------------------------
Number of entries          2
===============================================================================
*A:PE1#
```

The following command shows the addresses that are still free in a particular subnet:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" free-addresses 10.11.0.0/
24

===============================================================================
Free addresses
```

```
===============================================================================
IP Address      Fail Ctrl
-------------------------------------------------------------------------------
10.11.0.11      local
10.11.0.12      local
10.11.0.13      local
10.11.0.15      local
10.11.0.16      local
--- snipped ---
10.11.0.253     local
10.11.0.254     local
-------------------------------------------------------------------------------
No. of free addresses: 241
===============================================================================
*A:PE1#
```

The following command shows the DHCP server associations; this is the list of interfaces that the DHCP server is listening on:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" associations
===============================================================================
DHCP server dhcp4-VPRN1  router 1
===============================================================================
Interface associations
Interface                       Admin
-------------------------------------------------------------------------------
int-VPRN1-DHCPv4                Up

-------------------------------------------------------------------------------
Local Address Assignment associations
Group interface                 Admin
-------------------------------------------------------------------------------
===============================================================================
*A:PE1#
```

The following configuration enables debugging for DHCP server *dhcp4-VPRN1* on VPRN 1:

```
debug
    router "1"
        local-dhcp-server "dhcp4-VPRN1"
            detail-level high
            mode egr-ingr-and-dropped
        exit
    exit
exit
```

To ensure that the debug output is sent to a session, the following additional configuration is needed:

```
configure
    log
        log-id 1
            description "Send debug log to the current telnet/ssh session"
            from debug-trace
            to session
```

```
            no shutdown
        exit
    exit
exit
```

With this configuration, the following output is shown when the IPoE host with MAC address 00:00:00:01:01:01 connects:

```
13 2016/10/17 18:51:12.30 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
Rx DHCP Discover

   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 10.10.0.1
   chaddr: 00:00:00:01:01:01    xid: 0x21

   DHCP options:
   [82] Relay agent information: len = 23
      [1] Circuit-id: PE1|1|int-GRP|1/1/1:1
   [53] Message type: Discover
   [255] End

   Hex Packet Dump:
   01 01 06 00 00 00 00 21 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 0a
   --- snipped ---
   00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
"

14 2016/10/17 18:51:12.30 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
lease added for 10.10.0.12 state=offer
"

15 2016/10/17 18:51:12.30 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
Tx DHCP Offer to local relay agent 10.10.0.1 vrId=2

   ciaddr: 0.0.0.0           yiaddr: 10.10.0.12
   siaddr: 10.11.11.1        giaddr: 10.10.0.1
   chaddr: 00:00:00:01:01:01    xid: 0x21

   DHCP options:
   [82] Relay agent information: len = 23
      [1] Circuit-id: PE1|1|int-GRP|1/1/1:1
   [53] Message type: Offer
   [54] DHCP server addr: 10.11.11.1
   [51] Lease time: 600
   [1] Subnet mask: 255.255.255.0
   [3] Router: 10.10.0.1
   [130] Unknown option: len = 9, value = 4d 79 4f 70 74 69 6f 6e 31
   [6] Domain name server: length = 8
            1.1.1.1
            1.1.2.2
   [150] Unknown option: len = 4, value = 01 01 01 01
   [255] End

   Hex Packet Dump:
   02 01 06 00 00 00 00 21 00 00 00 00 00 00 00 00 0a 0a 00 0c 0a 0b 0b 01 0a
```

```
   --- snipped ---
   31 7c 31 7c 69 6e 74 2d 47 52 50 7c 31 2f 31 2f 31 3a 31 ff
"


16 2016/10/17 18:51:12.32 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
Rx DHCP Request

   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 10.10.0.1
   chaddr: 00:00:00:01:01:01    xid: 0x21

   DHCP options:
   [82] Relay agent information: len = 23
      [1] Circuit-id: PE1|1|int-GRP|1/1/1:1
   [53] Message type: Request
   [50] Requested IP addr: 10.10.0.12
   [54] DHCP server addr: 10.11.11.1
   [255] End

   Hex Packet Dump:
   01 01 06 00 00 00 00 21 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 0a
   --- snipped ---
   31 2f 31 3a 31 ff 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
"


17 2016/10/17 18:51:12.32 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
lease update for 10.10.0.12 state=stable
"


18 2016/10/17 18:51:12.52 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
Tx DHCP Ack to local relay agent 10.10.0.1 vrId=2

   ciaddr: 0.0.0.0           yiaddr: 10.10.0.12
   siaddr: 10.11.11.1        giaddr: 10.10.0.1
   chaddr: 00:00:00:01:01:01    xid: 0x21

   DHCP options:
   [82] Relay agent information: len = 23
      [1] Circuit-id: PE1|1|int-GRP|1/1/1:1
   [53] Message type: Ack
   [54] DHCP server addr: 10.11.11.1
   [51] Lease time: 600
   [1] Subnet mask: 255.255.255.0
   [3] Router: 10.10.0.1
   [130] Unknown option: len = 9, value = 4d 79 4f 70 74 69 6f 6e 31
   [6] Domain name server: length = 8
            1.1.1.1
            1.1.2.2
   [150] Unknown option: len = 4, value = 01 01 01 01
   [255] End

   Hex Packet Dump:
   02 01 06 00 00 00 00 21 00 00 00 00 00 00 00 00 0a 0a 00 0c 0a 0b 0b 01 0a
```

```
   --- snipped ---
   31 7c 31 7c 69 6e 74 2d 47 52 50 7c 31 2f 31 2f 31 3a 31 ff
"
```

When a client terminates its connection, the following output is shown:

```
19 2016/10/17 18:52:05.97 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
Rx DHCP Release

   ciaddr: 10.10.0.12        yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:00:00:01:01:01    xid: 0x21

   DHCP options:
   [82] Relay agent information: len = 23
      [1] Circuit-id: PE1|1|int-GRP|1/1/1:1
   [53] Message type: Release
   [54] DHCP server addr: 10.11.11.1
   [255] End

   Hex Packet Dump:
   01 01 06 00 00 00 00 21 00 00 00 00 0a 0a 00 0c 00 00 00 00 00 00 00 00
   --- snipped ---
   00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
"


20 2016/10/17 18:52:05.96 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
lease deleted for 10.10.0.12 (rxd release)
"
```

A PPP user connecting via local address assignment shows the following messages:

```
21 2016/10/17 18:52:15.97 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
Rx internal Request
   primary pool  : pool-2
   ciaddr        : 0.0.0.0
"

22 2016/10/17 18:52:15.97 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
lease added for 10.20.0.22 state=internal
"
```

When this user terminates the PPP session, the following messages are shown:

```
23 2016/10/17 18:52:26.41 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
Rx internal Release
   ciaddr        : 10.20.0.22
"

24 2016/10/17 18:52:26.41 CEST MINOR: DEBUG #2001 vprn1 DHCP server
```

```
"DHCP server:  dhcp4-VPRN1
lease deleted for 10.20.0.22 (rxd internal release)
"
```

Renumbering – Subnet Mask Change

The baseline configuration has the subnet 10.10.0.0/25 defined, providing address space for up to 126 addresses. The range that the server can take free addresses from starts at 10.10.0.11, and ends at 10.10.0.126, excluding the 10.10.0.61 to 10.10.70 sub-range.

Assume that subnet 10.10.0.128/25 was removed from a different BNG, and now can be used in this BNG. This subnet can be aggregated with the 10.10.0.0/25 network to become subnet 10.10.0.0/24. At the same time, the requirement is to not disrupt services for already connected users.

The following steps are required at the DHCP server:

- ensure that **no force-renews** is active
- delete the original subnet
- create the new subnet

Preventing the DHCP server from sending force-renew messages is important so that already connected users do not lose their connection, as follows:

```
*A:PE1# configure service vprn 1 dhcp local-dhcp-server dhcp4-VPRN1 no force-renews
```

The following command deletes the original subnet:

```
*A:PE1# configure service vprn 1 dhcp local-dhcp-server "dhcp4-VPRN1"
                                        pool "pool-1" no subnet 10.10.0.0/25
```

Leases are not deleted when the subnet is deleted; their status changes from *stable* to *removePending*, as follows:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" leases

===============================================================================
Leases for DHCP server dhcp4-VPRN1 router 1
===============================================================================
IP Address      Lease State       Mac Address       Remaining   Clnt  Fail
  PPP user name/Opt82 Circuit Id                    LifeTime    Type  Ctrl
  User-db/Sticky-lease Hostname
-------------------------------------------------------------------------------
10.10.0.11      removePending     00:00:00:01:01:01 1h57m25s    dhcp  local

10.10.0.12      removePending     00:00:00:03:01:01 0h57m30s    ppp   local
  PE1|1|int-GRP|1/1/1:1
```

```
-------------------------------------------------------------------------------
2 leases found
===============================================================================
*A:PE1#
```

This status change is also shown in the debug log, as follows:

```
132 2016/10/14 14:10:57.66 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
lease 10.10.0.11 scheduled for removal
"

133 2016/10/14 14:10:57.66 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
lease 10.10.0.12 scheduled for removal
"

134 2016/10/14 14:10:57.66 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
lease 10.10.0.11 scheduled for removal
"

135 2016/10/14 14:10:57.66 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
lease 10.10.0.12 scheduled for removal
"
```

Users trying to renew or connect will not get an address as long as no new subnet is defined.

Create the new 10.10.0.0/24 subnet, with the new address range starting at 10.10.0.11 and ending at 10.10.0.254, as follows. The original exclusion range still applies, but a new exclusion address 10.10.0.129 is added, to be described later:

```
configure
    service
        vprn 1 customer 1 create
            dhcp
                local-dhcp-server dhcp4-VPRN1 create
                    use-gi-address scope pool
                    no force-renews
                    pool "pool-1" create
                        options
                            dns-server 1.1.1.1 1.1.2.2
                            lease-time hrs 2
                            custom-option 150 address 1.1.1.1
                        exit
                        subnet 10.10.0.0/24 create
                            options
                                subnet-mask 255.255.255.0
                                default-router 10.10.0.1
                            exit
                            exclude-addresses 10.10.0.61 10.10.0.70
                            exclude-addresses 10.10.0.129 10.10.0.129
                            address-range 10.10.0.11 10.10.0.254
                        exit
```

```
                    exit
                exit
            exit
        exit
    exit
exit
```

Leases that were in use before return to the *stable* state, if they are not in the exclusion range, as follows:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" leases

===============================================================================
Leases for DHCP server dhcp4-VPRN1 router 1
===============================================================================
IP Address      Lease State       Mac Address       Remaining   Clnt  Fail
  PPP user name/Opt82 Circuit Id                    LifeTime    Type  Ctrl
  User-db/Sticky-lease Hostname
-------------------------------------------------------------------------------
10.10.0.11      stable            00:00:00:01:01:01 1h49m43s    dhcp  local

10.10.0.12      stable            00:00:00:03:01:01 0h49m48s    ppp   local
  PE1|1|int-GRP|1/1/1:1
-------------------------------------------------------------------------------
2 leases found
===============================================================================
*A:PE1#
```

The following command adds the 10.10.0.129/25 address to the *int-SUB* subscriber interface, so that offers in the 10.10.0.128/25 range will not get dropped by the relay agent. Any address in the 10.10.0.128/25 subnet could be used; the lowest one is used in this example. Because this address is in use by the subscriber interface, it must be added to the exclusion list in the DHCP server, as follows:

```
*A:PE1# configure service vprn 1 subscriber-interface int-SUB address 10.10.0.129/25
```

This configuration ensures service continuity for already connected subscribers. They will get their new /24 subnet when they renew or rebind their lease. No change is needed at the relay agent.

Merging the two subnets at the subscriber interface is only possible with a service interruption, because the subscriber interface addresses cannot be deleted when leases are in use. Also the Gi address configured in the dhcp context must be deleted.

```
*A:PE1>config>service>vprn>sub-if# no address 10.10.0.1/25
INFO: PIP #1398 Cannot delete/
change address when managed ARPs or leases defined for this subnet exist -
 1 managed-arps or leases exist
*A:PE1>config>service>vprn>sub-if#
```

To also merge the subnets at the subscriber interface, all the leases in these subnets must be deleted. When the address defined at the subscriber interface is also used as the Gi address by the relay agent, the Gi address must be removed first. Then, the subscriber interface address can be deleted and recreated with the correct netmask. Also, the Gi address can be redefined after that. The changes at the DHCP server are similar to the ones defined previously.

Renumbering – Subnet Migration

The following changes to the baseline configuration have to be made to support the migration of DHCP clients from the 10.10.0.0/25 and 10.11.0.0/24 subnets to the 10.12.0.0/20 subnet. For that purpose, the 10.10.0.0/25 and the 10.11.0.0/24 subnets have the keyword *drain* added, so that leases in the corresponding address ranges will not be extended.

This new 10.12.0.0/20 subnet has a new subnet mask, a new default router, and three address ranges. New clients connecting will automatically get addresses from this new subnet. To ensure existing clients will not lose their connection, the **use-gi-address scope** is set to **pool**, so that they get a new lease from the new subnet when renewing or rebinding.

In scenarios where lease times are long (an order of magnitude of months or even years), it can take a considerable time before all clients have a lease in the new subnet. Having DHCP clients supporting force-renew can help speed up the migration process. The following configuration has force-renews enabled.

Address 10.12.0.1 is used as the default router for this subnet, so this address is added to the *int-SUB* subscriber-interface. This address will later be used as the Gi address.

```
configure
    service
        vprn 1
            dhcp
                local-dhcp-server "dhcp4-VPRN1" create
                    use-gi-address scope pool
                    force-renews
                    pool "pool-1" create
                        options
                            dns-server 1.1.1.1 1.1.2.2
                            lease-time hrs 2
                        exit
                        subnet 10.10.0.0/25 create
                            drain
                            options
                                subnet-mask 255.255.255.0
                                default-router 10.10.0.1
                            exit
```

```
                            address-range 10.10.0.11 10.10.0.12
                        exit
                        subnet 10.11.0.0/24 create
                            drain
                            options
                                subnet-mask 255.255.255.0
                                default-router 10.10.0.1
                            exit
                            address-range 10.11.0.11 10.11.0.254
                        exit
                        subnet 10.12.0.0/20 create
                            options
                                subnet-mask 255.255.240.0
                                default-router 10.12.0.1
                            exit
                            address-range 10.12.0.10 10.12.12.255
                            address-range 10.12.13.1 10.12.14.255
                            address-range 10.12.15.10 10.12.15.254
                        exit
                    exit
                    no shutdown
                exit
            exit
            subscriber-interface "int-SUB"
                address 10.12.0.1/20
            exit
        exit
    exit
exit
```

The following command shows that the original subnets are in the drained state:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" summary
===============================================================================
DHCP server dhcp4-VPRN1  router 1
===============================================================================
Admin State            : inService
Operational State      : inService
Persistency State      : ok

--- snipped ---

-------------------------------------------------------------------------------
Pool name : pool-1
-------------------------------------------------------------------------------
Failover Admin State   : outOfService
Failover Oper State    : shutdown
Failover Persist Key   : 0x00000001
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
-------------------------------------------------------------------------------
10.10.0.0/25           2        100% 0        0        0        0        Y
10.11.0.0/24           244      100% 0        0        0        0        Y
```

```
10.12.0.0/20           4072     99%  3        0        0        0        N
Totals for pool        4318     99%  3        0        0        0
-------------------------------------------------------------------------------

Totals for server      4318     99%  3        0        0        0

--- snipped ---

===============================================================================
*A:PE1#
```

Because the DHCP server is configured with force-renew, connected clients are sent a force-renew message. In response, the client tries to extend its lease by sending a request message using the current address. The DHCP server sends a negative-acknowledgement (NAK) to the requesting client, because the subnet is in the draining state. This forces the client to go through the DHCP initialization process; a new DORA message sequence is initiated. Therefore, the client gets a free address in the new subnet, with a new netmask, and a new default router, as follows. The same DNS servers are offered, because these pool options were not changed.

```
1 2016/10/15 19:19:36.04 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
Tx DHCP ForceRenew to client at 10.10.0.12 vrId=2

   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 10.11.11.1        giaddr: 0.0.0.0
   chaddr: 00:00:00:01:01:01    xid: 0x1f

   DHCP options:
   [53] Message type: ForceRenew
   [54] DHCP server addr: 10.11.11.1
   [255] End
"

2 2016/10/15 19:19:36.05 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
Rx DHCP Request

   ciaddr: 10.10.0.12        yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:00:00:01:01:01    xid: 0x1f

   DHCP options:
   [53] Message type: Request
   [255] End
"

3 2016/10/15 19:19:36.05 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
subnet is draining
Tx DHCP Nak to client 10.10.0.12 vrId=2 (via snooping function)

   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:00:00:01:01:01    xid: 0x1f
```

```
   DHCP options:
   [53] Message type: Nak
   [54] DHCP server addr: 10.11.11.1
   [255] End
"

4 2016/10/15 19:19:36.06 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
Rx DHCP Discover

   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 10.10.0.1
   chaddr: 00:00:00:01:01:01    xid: 0x1f

   DHCP options:
   [53] Message type: Discover
   [255] End
"

5 2016/10/15 19:19:36.06 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
lease added for 10.12.0.17 state=offer
"

6 2016/10/15 19:19:36.06 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
Tx DHCP Offer to local relay agent 10.10.0.1 vrId=2

   ciaddr: 0.0.0.0           yiaddr: 10.12.0.17
   siaddr: 10.11.11.1        giaddr: 10.10.0.1
   chaddr: 00:00:00:01:01:01    xid: 0x1f

   DHCP options:
   [53] Message type: Offer
   [54] DHCP server addr: 10.11.11.1
   [51] Lease time: 7200
   [1] Subnet mask: 255.255.240.0
   [3] Router: 10.12.0.1
   [6] Domain name server: length = 8
            1.1.1.1
            1.1.2.2
   [150] Unknown option: len = 4, value = 01 01 01 01
   [255] End
"

7 2016/10/15 19:19:36.07 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
Rx DHCP Request

   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 10.10.0.1
   chaddr: 00:00:00:01:01:01    xid: 0x1f

   DHCP options:
   [53] Message type: Request
   [50] Requested IP addr: 10.12.0.17
   [54] DHCP server addr: 10.11.11.1
   [255] End
"
```

```
8 2016/10/15 19:19:36.07 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
lease update for 10.12.0.17 state=stable
"

9 2016/10/15 19:19:36.24 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
Tx DHCP Ack to local relay agent 10.10.0.1 vrId=2

   ciaddr: 0.0.0.0           yiaddr: 10.12.0.17
   siaddr: 10.11.11.1        giaddr: 10.10.0.1
   chaddr: 00:00:00:01:01:01    xid: 0x1f

   DHCP options:
   [53] Message type: Ack
   [54] DHCP server addr: 10.11.11.1
   [51] Lease time: 7200
   [1] Subnet mask: 255.255.240.0
   [3] Router: 10.12.0.1
   [6] Domain name server: length = 8
            1.1.1.1
            1.1.2.2
   [150] Unknown option: len = 4, value = 01 01 01 01
   [255] End
"
```

When the original DHCP server subnets are fully drained, they can be safely deleted. However, deleting a subnet from a pool before it is fully drained results in the remaining leases being scheduled for removal, as follows:

```
140 2016/10/10 15:12:42.87 CEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  dhcp4-VPRN1
lease 10.11.0.11 scheduled for removal
"
```

The number of leases pending for removal can be shown using following command:

```
*A:PE1# show router 1 dhcp local-dhcp-server "dhcp4-VPRN1" summary
===============================================================================
DHCP server dhcp4-VPRN1  router 1
===============================================================================
Admin State            : inService
Operational State      : inService

--- snipped ---

-------------------------------------------------------------------------------
Pool name : pool-1
-------------------------------------------------------------------------------
Failover Admin State   : outOfService
Failover Oper State    : shutdown
Failover Persist Key   : 0x00000001
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
-------------------------------------------------------------------------------
10.11.0.0/24           244      100% 0        0        0        0        Y
10.12.0.0/20           4075     100% 0        0        0        0        N
Totals for pool        4319     100% 0        0        0        0
-------------------------------------------------------------------------------
Not subnet related                                                 Rem-pend
-------------------------------------------------------------------------------
                                                                   1
-------------------------------------------------------------------------------

Totals for server      4319     100% 0        0        0        1

--- snipped ---

===============================================================================
*A:PE1#
```

This lease will be deleted when the lease expires.

The relay agent can then have the Gi address updated (10.12.0.1) and the old subnets can be removed from the group interface.

Conclusion

SR OS supports DHCPv4 servers on any routing instance (VPRN or base router), offering pool, subnet, and address management, combined with configuration parameter management and persistency.

Diameter Application NASREQ

This chapter provides information about Diameter Application NASREQ.

Topics in this chapter include:

Applicability

This configuration note is applicable to the 7750 SR. The configuration was tested on release 13.0.R3.

Overview

NASREQ is defined in RFC 7155, *Diameter Network Access Server Application*, (obsoletes RFC 4005). Because NASREQ is a Diameter application, it uses the Diameter base protocol defined in RFC 6733, *Diameter Base Protocol*. The purpose of NASREQ in 7750 SR Release 13.0 is to provide subscriber authentication and authorization. NASREQ provides functionality that is also available for RADIUS, but uses the Diameter protocol instead.

NASREQ complements the other Diameter applications supported in the SR OS:

- Gx provides advanced authorization capabilities for subscribers and usage monitoring, and interfaces with a Policy and Charging Rules Function (PCRF).
- Gy or Diameter Credit Control Application (DCCA) provides on-line charging functionality, and interfaces with an On-line Charging Server (OCS).
- NASREQ provides subscriber authentication and authorization, and interfaces with an AAA server.

These three Diameter applications use the Diameter base protocol, which is described in the Establishing a Diameter Peering Session chapter of this Advanced Configuration Guide. That chapter also describes the configuration of the Diameter base protocol, which is very similar for all supported Diameter applications, and is not repeated in this chapter.

When a subscriber connects to the BNG, NASREQ is triggered, as would have been the case with RADIUS authentication. The following figure shows a sample call flow for NASREQ and Gx applications when an IPoE/IPv4 subscriber connects. The supported NASREQ messages are Authentication and Authorization Request (AAR) and Authentication and Authorization Answer (AAA), and the BNG assumes that the AAA server does not maintain session state. Therefore, there is no need for the BNG to send a message to the AAA server to indicate when the subscriber has ended the session; for example, by sending a DHCP release. The BNG negotiates with the AAA server that it expects a stateless behavior by sending the Auth-Session-State AVP with value NO_STATE_MAINTAINED (1) as defined in RFC 6733, and the AAA server has to confirm that by sending back the AVP with the same value.

Figure 92 **NASREQ Trigger**

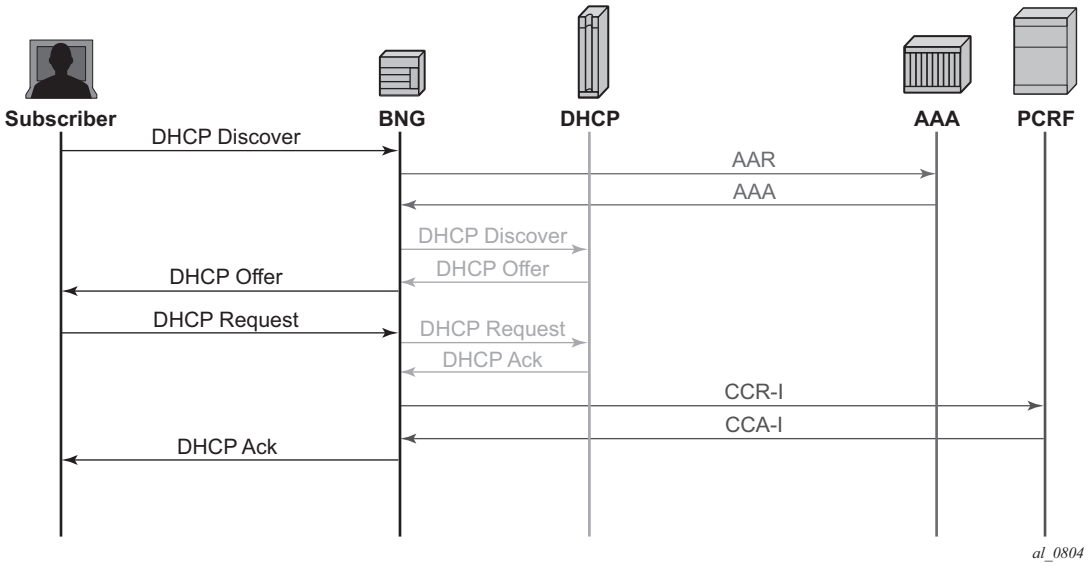

As shown in Figure 92, NASREQ is triggered when the DHCP discover is received, while Gx, using the Credit Control Request (CCR-I) and Credit Control Answer (CCA-I) messages, is triggered at the end, after the IP address allocation.

After the subscriber is authenticated, no re-authentication is required, except for IPoE subscribers, when there is a DHCP renew and the relay agent information changes.

Supported NASREQ AVPs in AAA are listed in the **7750 SR Gx AVPs Reference Guide**, and include the following AVPs:

- Standard NASREQ authorization AVPs supported:

Table 14 **Standard NASREQ authorization AVPs**

AVP #	AVP Name
7	Framed-Protocol
8	Framed-IP-Address
9	Framed-IP-Netmask
22	Framed-Route
25	Class
88	Framed-Pool
97	Framed-IPv6-Prefix
99	Framed-IPv6-Route
100	Framed-IPv6-Pool

- Vendor-specific NASREQ authorization AVPs supported:

Table 15 **Vendor-specific NASREQ authorization AVPs**

AVP #	AVP Name		AVP #	AVP Name
v-6527-9	Alc-Primary-DNS		v-6527-33	Alc-MSAP-Interface
v-6527-10	Alc-Secondary-Dns		v-6527-45	Alc-App-Prof-Str (**Gx)**
v-6527-11	Alc-Subsc-ID-Str		v-6527-99	Alc-Ipv6-Address
v-6527-12	Alc-Subsc-Prof-Str (**Gx)**		v-6527-105	Alc-Ipv6-Primary-Dns
v-6527-13	Alc-SLA-Prof-Str (**Gx)**		v-6527-106	Alc-Ipv6-Secondary-Dns
v-6527-16	Alc-ANCP-Str		v-6527-131	Alc-Delegated-IPv6-Pool
v-6527-17	Alc-Retail-Serv-Id		v-6527-161	Alc-Delegated-IPv6-Prefix-Length
v-6527-18	Alc-Default-Router		v-6527-174	Alc-Lease-Time
v-6527-28	Alc-Int-Dest-Id-Str (**Gx)**		v-6527-181	Alc-SLAAC-IPv6-Pool
v-6527-29	Alc-Primary-Nbns			

Table 15 **Vendor-specific NASREQ authorization AVPs (Continued)**

AVP #	AVP Name		AVP #	AVP Name
v-6527-30	Alc-Secondary-Nbns			
v-6527-31	Alc-MSAP-Serv-Id			
v-6527-32	Alc-MSAP-Policy			

All the NASREQ AVPs are also supported by RADIUS, and their meaning is described in the RADIUS AVP Reference Guide. Most of the NASREQ AVPs are related to IP address assignment to the subscriber and service selection in the case of managed SAPs and/or wholesale/retail. These AVPs are not available in Gx, because Gx is triggered after service selection and IP address assignment. The four AVPs marked with **Gx** in the preceding list are supported by both NASREQ and Gx. If both the AAA server and PCRF return the same attribute, but with a different value, then the value from the PCRF will be used to create the subscriber.

The AAA server authenticates a subscriber based on a password. For DHCP, the password is configured in the BNG because DHCP cannot provide a subscriber password, and for PPP, the authentication is performed through PAP or CHAP.

Configuration

Configuration of NASREQ is performed in four steps:

Step 1. Configure a Diameter peer policy.

Step 2. Configure a Diameter application policy.

Step 3. Assign the Diameter application policy.

Step 4. Apply a Python policy (optional).

In the first step, the Diameter peer policy configures the Diameter base protocol and is described in the Establishing a Diameter Peering Session chapter. Note that it is possible to enable both Gx and NASREQ at the same time in one Diameter peer policy:

```
configure aaa diameter-peer-policy "DSC.NASREQ.Gx.39.122"
     applications gx nasreq
```

In this example, the same Diameter connection will be used for both Gx and NASREQ.

In the second step, the Diameter application policy configures the NASREQ application. An example configuration of a Diameter application policy for NASREQ follows:

```
configure subscriber-mgmt diameter-application-policy "DSC.NASREQ"
     description "Diameter application policy for Nasreq to DSC 39.122"
     application nasreq
     diameter-peer-policy "DSC.NASREQ.Gx.39.122"
     nasreq
          password "5bBaQIsc4It1Ji/WjmPu2oVb2Bg883Om3ex1mCs8xYE" hash2
          user-name-format circuit-id
          user-name-operation append-domain domain "MyDomain"
          include-avp
               circuit-id
               nas-port-id prefix-type user-string prefix-string "MyBNG"\
                           suffix-type user-string suffix-string "Belgium"
               nas-port-type
               remote-id
          exit
      exit
```

The application nasreq has to be specified. Contrary to the Diameter peer policy, only one application can be enabled in the Diameter application policy. If both Gx and NASREQ need to use the same Diameter peer policy, two Diameter application policies have to be configured, both referring to the same Diameter peer policy with the command diameter-peer-policy.

The password command configures the password to be used for IPoE subscribers toward the AAA server. For PPP subscribers, the subscriber credentials will come from the PAP or CHAP authentication.

The AAA server needs the identity of the subscriber, which is in the User-Name AVP. The username can be configured with the command user-name-format. Options for the username format are, for example:

- MAC address (with or without giaddr),
- circuit-id from the relay agent information (for example, DHCP Option 82 for IPv4),
- information from DHCP option 60 and 61 (which contain the Client-id and Vendor-Class information),
- and NAS port Id.

The username can be further formatted with the command user-name-operation, which applies to both IPoE and PPP subscribers. This command allows you to add, remove, or replace a domain name. The domain name is defined as the part of the username that follows the first @ symbol in the username; for example, if the username is MyName@MyISP.com, then the domain name in this example is MyISP.com. If a domain name is added, and there is already a domain name, then

the second domain name will be added after the first one, and a dot (.) will be used as separator between the 2 domain names. For example, with the preceding configuration, if the circuit-id is 1/1/1, then the resulting username will be 1/1/1@MyDomain, and if the circuit-id is 1/1/1@DSLAM1, then the username will be 1/1/1@DSLAM1.MyDomain.

The following extra AVPs can be included with the command include-avp:

- circuit-id and remote-id from the relay agent information (for example, Option 82 in DHCP),
- nas-port,
- nas-port-id,
- called-station-id,
- calling-station-id.

The nas-port-id can include a prefix string and a suffix string or circuit-id/remote-id as suffix. The value for nas-port-type can be specified, but if nas-port-type is enabled without a value, then it will be derived from the type of the actual nas-port-id. In the preceding example, the nas-port-id is prepended with the string MyBNG and the string Belgium is appended to the nas-port-id. If the nas-port-id is, for example, 1/1/4:3001 (that is the subscriber is connecting to SAP 1/1/4:3001), then the eventual nas-port-id value with the preceding configuration will be MyBNG 1/1/4:3001Belgium and the nas-port-type will be 15, indicating Ethernet.

If an AVP contains a MAC address, the format of the MAC address can be specified as follows:

```
configure subscriber-mgmt diameter-application-policy \
                                        "DSC.NASREQ" nasreq mac-format
  - mac-format <mac-format>
  - no mac-format

 <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
```

The format will be applied to all AVPs containing MAC addresses; for example, User-Name and Calling-Station-id AVPs.

If more than one AAA server is used for redundancy, the Diameter peer policy has to be configured with a peer per AAA server, and in the Diameter application policy, the failure handling has to be configured. Recommended configuration options for failure handling are:

- if there is one AAA, and new sessions have to be accepted when the AAA is down, then configure the Diameter application policy as follows:

```
on-failure failover disabled handling continue
```

- if there is one AAA, and new sessions have to be rejected when the AAA is down, then configure the Diameter application policy as follows:

```
on-failure failover disabled handling terminate
```

- if there are two or more AAA servers, configure the Diameter application policy as follows:

```
on-failure failover enabled handling retry-and-terminate
```

The TX-timer for NASREQ is also configurable and is 10 seconds by default. The TX-timer is the time that the BNG waits to get an answer from the AAA server before applying the configured failure handling.

In the third and last mandatory step, the Diameter application policy has to be applied using the command diameter-auth-policy to any of the following:

- group interface in a VPRN or IES
- local user database (LUDB)
- capture SAP

The fourth configuration step is optional and allows a Python script to modify the AAR and AAA messages. A sample script that copies the User-Name AVP into the Subscription-Id AVP is as follows:

```
file type cf3:/dsc_subscriber_id
File: dsc_subscriber_id
------------------------------------------------------------------------
from binascii import *
from alc import diameter

def getint_b2a_hex(val):
    return  int(b2a_hex(val),16)

def byte_to_binary(n):
    return ''.join(str((n & (1 << i)) and 1) for i in reversed(range(8)))

def hex_to_binary(h):
    return ''.join(byte_to_binary(ord(b)) for b in unhexlify(h))

def checkbitset(byte,index):
    return ((byte&(1<<index))!=0)

def checkRequestOrReply():
    if checkbitset(int(hex_to_binary(b2a_hex(diameter.flags))),7) is True:
        return 'R'
    else :
        return 'A'

# prints only for debug purpose
# print "Script Name - subscr_avp_modify.py add Subscription-ID"
```

```
if getint_b2a_hex(diameter.code) == 265 and checkRequestOrReply() == 'R':
    try:
         username = diameter.get_avps(1,0)[0][1]
         if username != "":
# prints only for debug purpose
#          print "username is " + str(username)
           diameter.set_grouped_avps(443,0,[('@', {(450,0): [('@','\x00\x00\x00\x04'
)], (444,0): [('@',str(username))]})])
    except Exception, err: print "Python FAILED to fetch username"

========================================================================
```

The preceding Python script can be configured in a Python policy as follows:

```
configure python python-script "dsc_subscriber_id" create
            primary-url "cf3:/dsc_subscriber_id"
            no shutdown
        exit

configure python python-policy "dsc_subscriber_id" create
            description "add Subscription-Id to NASREQ AAR"
            diameter aar direction egress script "dsc_subscriber_id"
        exit
```

The Python policy applies the script to the AAR messages in the egress direction. In a similar way, Python scripts can also be applied to AAA messages in the ingress direction.

After the Python policy has been configured, it has to be applied in the Diameter peer policy:

```
configure aaa diameter-peer-policy "DSC.NASREQ.Gx.39.122" \
   python-policy "dsc_subscriber_id"
```

With this configuration, all Diameter application policies that are using this Diameter peer policy will have their AAR messages (in the egress direction) modified, according to the Python script.

Troubleshooting

The operational state of the Diameter application policy can be shown with the following command. The command shows configuration details, as well as where the policy is applied. In the output, the policy is applied to the group interface **to_STC** in VPRN 10001 and to the capture SAP 1/1/5:* in VPLS 20002.

```
show subscriber-mgmt diameter-application-policy "DSC.NASREQ"
========================================================================
DIAMETER application policy "DSC.NASREQ"
```

```
===============================================================================
Description                : Diameter application policy for Nasreq to DSC 39.
                             122
Session failover           : enabled
Failover handling          : terminate
Peer policy                : DSC.NASREQ.Gx.39.122
Application                : nasreq
Tx timer (s)               : 10
Last management change     : 06/18/2015 12:58:47
-------------------------------------------------------------------------------
NASREQ
-------------------------------------------------------------------------------
Include AVP                : circuit-id
                             remote-id
                             called-station-id
                             calling-station-id
                             nas-port-id
                             nas-port-type
Calling-Station-Id type    : sap-string
NAS-Port-Id prefix type    : user-string
NAS-Port-Id prefix         : MyBNG
NAS-Port-Id suffix type    : user-string
NAS-Port-Id suffix         : Belgium
NAS-Port-Type type         : standard

User name format           : circuit-id
User name operation        : append-domain
Domain name                : MyDomain
MAC address format         : aa:
Last management change     : 06/18/2015 12:58:56
===============================================================================
No interfaces found using diameter-application-policy "DSC.NASREQ".
===============================================================================
Interfaces using diameter-auth-policy "DSC.NASREQ"
===============================================================================
Interface-name                   Service-id Type
-------------------------------------------------------------------------------
to_STC                           10001      VPRN
-------------------------------------------------------------------------------
No. of interfaces: 1
===============================================================================
-------------------------------------------------------------------------------
VPLS SAP's with diameter-auth-policy "DSC.NASREQ"
-------------------------------------------------------------------------------
Service    SAP
-------------------------------------------------------------------------------
20002      1/1/5:*
-------------------------------------------------------------------------------
No. of SAP's: 1
-------------------------------------------------------------------------------
```

Debugging of NASREQ is performed at the layer of the Diameter base protocol where the NASREQ specific messages have to be enabled. For example, to debug NASREQ, a debug configuration could be as follows:

```
debug
    diameter
        detail-level high
```

```
        no dest-realm
        no diameter-peer
        no diameter-peer-policy
        message-type aar aaa
        no origin-realm
    exit
exit
```

The **message-type aar** and **aaa** are specific for NASREQ. If a Python script is used, the Python script could be debugged as follows:

```
debug
    python
        python-script "dsc_subscriber_id"
            script-all-info
        exit
    exit
exit
```

Following is a debug output example of a DHCP Discover, AAR, and AAA for the previous configuration example.

```
69 2015/06/19 14:47:33.69 CET MINOR: DEBUG #2001 vprn10001 PIP
"PIP: DHCP
instance 16 (10001), interface index 3 (to_STC),
   received DHCP Boot Request on Interface to_STC (1/1/4:3001) Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:10:95:00:00:01    xid: 0x11

   DHCP options:
   [82] Relay agent information: len = 23
      [1] Circuit-id: 1/1/1
      [2] Remote-id: stefaansremote
   [53] Message type: Discover
   [57] Max msg size: 576
   [61] Client id: (hex) 01 00 10 95 00 00 01
   [51] Lease time: 60
   [12] Host name: client_Port //5/11-0-0
   [55] Param request list: len = 5
             1  Subnet mask
             6  Domain name server
            15  Domain name
            33  Static route
            44  NETBIOS name server
   [255] End

   Hex Packet Dump:
   01 01 06 00 00 00 00 11 00 00 80 00 00 00 00 00 00 00 00 00 00 00 00 00
   00 00 00 00 10 95 00 00 01 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 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 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 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 00 63 82 53 63 35 01 01 39 02 02 40 3d 07 01
   00 10 95 00 00 01 33 04 00 00 00 3c 0c 16 63 6c 69 65 6e 74 5f 50 6f 72 74
   20 2f 2f 35 2f 31 31 2d 30 2d 30 37 05 01 06 0f 21 2c 52 17 01 05 31 2f 31
   2f 31 02 0e 73 74 65 66 61 61 6e 73 72 65 6d 6f 74 65 ff 00
"

70 2015/06/19 14:47:33.68 CET MINOR: DEBUG #2001 Base Python Output
"Python Output: dsc_subscriber_id
"

71 2015/06/19 14:47:33.69 CET MINOR: DEBUG #2001 Base Python Result
"Python Result: dsc_subscriber_id
Diameter AVP code 443, SET
        "('@', '\\x00\\x00\\x01\\xc2@\\x00\\x00\\x0c\\x00\\x00\\x00\\x04\\x00\\x00\\
x01
\\xbc@\\x00\\x00\\x161/1/1@MyDomain\\x00\\x00')"
"

72 2015/06/19 14:47:33.69 CET MINOR: DEBUG #2001 vprn10000 DIAMETER
"DIAMETER: Message Transmission
AAR from [DSC.NASREQ.Gx.39.122, DSC.39.122] to 10.40.15.4:3868
  Header
    ver 1  len 340  flags RP------  code 265
    app-id 1  hbh-id 15215  e2e-id 3204428864
  AVPs
    session-Id (263) -M------ [39]
      data [31] (UTF8String) : wlangw-2.SRrealm;1434532511;259
    auth-appl-id (258) -M------ [12]
      data [4] (Unsigned32) : 1 : Nasreq
    origin-host (264) -M------ [24]
      data [16] (DiameterIdentity) : wlangw-2.SRrealm
    origin-realm (296) -M------ [15]
      data [7] (DiameterIdentity) : SRrealm
    destination-realm (283) -M------ [16]
      data [8] (DiameterIdentity) : dscrealm
    auth-request-type (274) -M------ [12]
      data [4] (Enumerated) : 3 : AUTHORIZE_AUTHENTICATE
    nas-port-id (87) -M------ [31]
      data [23] (UTF8String) : MyBNG 1/1/4:3001Belgium
    nas-port-type (61) -M------ [12]
      data [4] (Enumerated) : 15 : Ethernet
    origin-state-id (278) -M------ [12]
      data [4] (Unsigned32) : 1434532511
    user-name (1) -M------ [22]
      data [14] (UTF8String) : 1/1/1@MyDomain
    user-password (2) -M------ [16]
      data [8] (OctetString) : 0x69 70 74 61 63 31 32 33
    auth-session-state (277) -M------ [12]
      data [4] (Enumerated) : 1 : NO_STATE_MAINTAINED
    agent-circuit-id (1) VM------ [17]
      vendor-id DSL_FORUM
      data [5] (OctetString) : 0x31 2f 31 2f 31
    agent-remote-id (2) VM------ [26]
      vendor-id DSL_FORUM
      data [14] (OctetString) : 0x73 74 65 66 61 61 6e 73 72 65 6d 6f 74 65
    subscription-id (443) -M------ [44]
      data [36] (Grouped)
```

```
        subscription-id-type (450) -M------ [12]
          data [4] (Enumerated) : 4 : private
        subscription-id-data (444) -M------ [22]
          data [14] (UTF8String) : 1/1/1@MyDomain

  01 00 01 54  c0 00 01 09  00 00 00 01  00 00 3b 6f
  be ff b4 40  00 00 01 07  40 00 00 27  77 6c 61 6e
  67 77 2d 32  2e 53 52 72  65 61 6c 6d  3b 31 34 33
  34 35 33 32  35 31 31 3b  32 35 39 00  00 00 01 02
  40 00 00 0c  00 00 00 01  00 00 01 08  40 00 00 18
  77 6c 61 6e  67 77 2d 32  2e 53 52 72  65 61 6c 6d
  00 00 01 28  40 00 00 0f  53 52 72 65  61 6c 6d 00
  00 00 01 1b  40 00 00 10  64 73 63 72  65 61 6c 6d
  00 00 01 12  40 00 00 0c  00 00 00 03  00 00 00 57
  40 00 00 1f  4d 79 42 4e  47 20 31 2f  31 2f 34 3a
  33 30 30 31  42 65 6c 67  69 75 6d 00  00 00 00 3d
  40 00 00 0c  00 00 00 0f  00 00 01 16  40 00 00 0c
  55 81 3a 9f  00 00 00 01  40 00 00 16  31 2f 31 2f
  31 40 4d 79  44 6f 6d 61  69 6e 00 00  00 00 00 02
  40 00 00 10  69 70 74 61  63 31 32 33  00 00 01 15
  40 00 00 0c  00 00 00 01  00 00 00 01  c0 00 00 11
  00 00 0d e9  31 2f 31 2f  31 00 00 00  00 00 00 02
  c0 00 00 1a  00 00 0d e9  73 74 65 66  61 61 6e 73
  72 65 6d 6f  74 65 00 00  00 00 01 bb  40 00 00 2c
  00 00 01 c2  40 00 00 0c  00 00 00 04  00 00 01 bc
  40 00 00 16  31 2f 31 2f  31 40 4d 79  44 6f 6d 61
  69 6e 00 00
"

73 2015/06/19 14:47:33.70 CET MINOR: DEBUG #2001 vprn10000 DIAMETER
"DIAMETER: Message Reception
AAA from 10.40.15.4:3868 to [DSC.NASREQ.Gx.39.122, DSC.39.122]
  Header
    ver 1  len 160  flags -P------  code 265
    app-id 1  hbh-id 15215  e2e-id 3204428864
  AVPs
    session-Id (263) -M------ [39]
      data [31] (UTF8String) : wlangw-2.SRrealm;1434532511;259
    origin-host (264) -M------ [22]
      data [14] (DiameterIdentity) : vm122.dscrealm
    origin-realm (296) -M------ [16]
      data [8] (DiameterIdentity) : dscrealm
    auth-appl-id (258) -M------ [12]
      data [4] (Unsigned32) : 1 : Nasreq
    result-code (268) -M------ [12]
      data [4] (Unsigned32) : 2001 : DIAM_RESCODE_SUCCESS
    auth-request-type (274) -M------ [12]
      data [4] (Enumerated) : 3 : AUTHORIZE_AUTHENTICATE
    auth-session-state (277) -M------ [12]
      data [4] (Enumerated) : 1 : NO_STATE_MAINTAINED
    origin-state-id (278) -M------ [12]
      data [4] (Unsigned32) : 1426768161

  01 00 00 a0  40 00 01 09  00 00 00 01  00 00 3b 6f
  be ff b4 40  00 00 01 07  40 00 00 27  77 6c 61 6e
  67 77 2d 32  2e 53 52 72  65 61 6c 6d  3b 31 34 33
  34 35 33 32  35 31 31 3b  32 35 39 00  00 00 01 08
  40 00 00 16  76 6d 31 32  32 2e 64 73  63 72 65 61
  6c 6d 00 00  00 00 01 28  40 00 00 10  64 73 63 72
```

```
  65 61 6c 6d  00 00 01 02  40 00 00 0c  00 00 00 01
  00 00 01 0c  40 00 00 0c  00 00 07 d1  00 00 01 12
  40 00 00 0c  00 00 00 03  00 00 01 15  40 00 00 0c
  00 00 00 01  00 00 01 16  40 00 00 0c  55 0a c1 21
"
```

Note that in the preceding example, the AAR contains a subscription-id AVP. This AVP is not included by default by the BNG, and also cannot be included by CLI, but it is added by the Python script dsc_subscriber_id.

Conclusion

This chapter describes NASREQ that provides subscriber authentication and authorization. NASREQ is based on Diameter, which is an evolution of RADIUS.

Diameter Inter-Chassis Redundancy

This chapter provides information about Diameter inter-chassis redundancy.

Topics in this chapter include:

Applicability

This chapter is applicable to 7750 SR.

The configuration was tested on release 13.0.R3.

Overview

A multi-homed broadband network gateway (BNG) setup consists of two physical BNGs, with one BNG providing redundancy for the other BNG (providing inter-chassis redundancy). To support Diameter for a multi-homed BNG setup, two solutions are possible:

1. Each physical BNG has a unique Diameter hostname/realm and establishes Diameter connections to its peers. This means that both physical BNGs have their Diameter connections up at the same time. This solution does not make use of the Diameter proxy supported by SR OS.
2. Using the SR OS Diameter proxy, both physical BNGs share the same Diameter hostname/realm, and only one physical BNG establishes a Diameter connection to its peers. In this way, the server only sees one logical BNG. The Diameter proxy is supported for the Gx and NASREQ Diameter applications.

With stateful Diameter applications, multiple messages are sent during a Diameter session. For instance, in the case of Gx, a single session has at least one credit control request-initial (CCR-I) message and one CCR-termination (CCR-T) message, and possibly also CCR-update (CCR-Us) messages. With the first solution above, when a BNG switchover occurs during the session, the CCR-I and CCR-T

messages are sent by different physical BNGs with different origin-host/realms. This might not be accepted by some policy and charging rules functions (PCRFs) that do not support a change of origin-host/realm during a session. With the second solution, all messages belonging to the same Diameter session will have the same origin-host/realm, irrespective of which BNG is active. However, having two devices deliberately configured with the same Diameter hostname and realm might confuse servers and Diameter Routing Agents (DRAs) as to where to send the Diameter messages.

With stateless Diameter applications, both solutions are suitable because only one message is sent during the Diameter session. There cannot be a change of origin-host/realm for a session when there is a BNG switchover.

The second solution, based on Diameter proxy, consists of three components as shown in Figure 93:

- ESM
- Diameter client
- Diameter proxy

Figure 93 **Diameter Proxy Implementation**

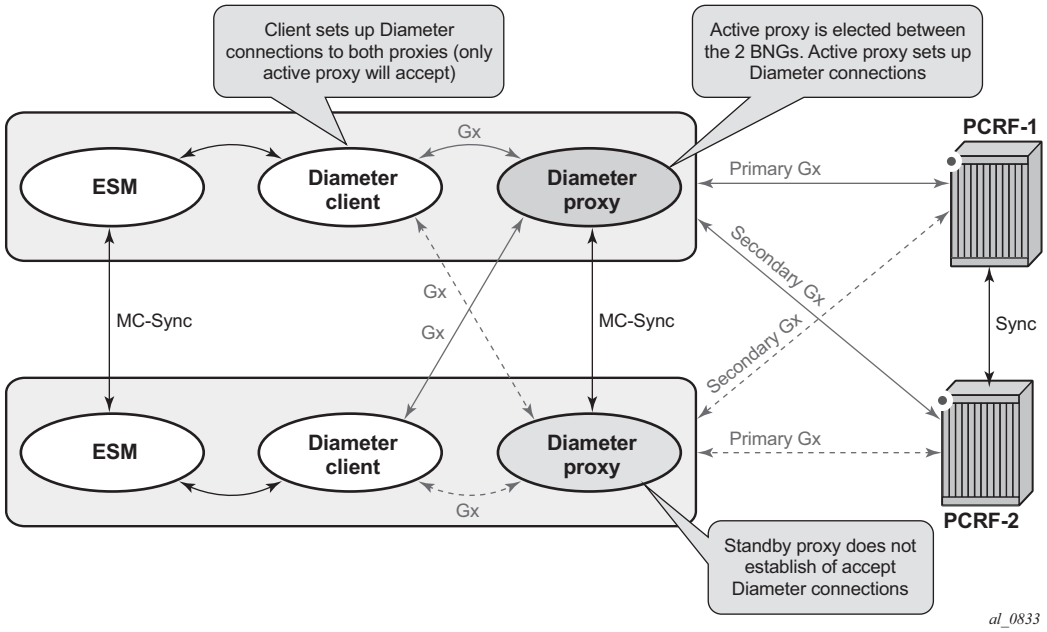

Figure 93 shows two physical BNGs; each of them have a Diameter client and proxy configured. The Diameter client will establish Diameter connections to the proxy, and in turn the proxy will establish Diameter connections to its peers in the network (PCRF, authorization, authentication and accounting (AAA), DRA, and so on).

An active and standby proxy will be selected, based on the highest chassis MAC address in a non-revertive way. When a proxy boots up, it first checks whether there is another proxy that is active. If it does not detect another active proxy within 10 min, this proxy will be selected as the active one. Only the active proxy will establish Diameter connections to the network, while the standby proxy will not establish any Diameter connections. This is indicated by the dotted lines in Figure 93.

The Diameter clients are configured with two Diameter peers, one to each proxy, and will try to establish both Diameter connections. However, only the active Diameter proxy will accept incoming Diameter connections. The standby proxy will reject any incoming Diameter connections by responding with a TCP-RST each time a client tries to connect.

The Diameter connection from the client to the proxy on the same physical BNG can be an internal connection, that is, it does not leave the BNG but it is a Diameter connection with the same functionality as an external connection.

When a subscriber connects to a Gx enabled BNG, a CCR-I is sent to the PCRF, see Figure 94. If the ESM master and the active proxy are on the same BNG, the CCR-I is sent from the client to the proxy on the same physical BNG:

1. A subscriber connects to the network; for example, by sending a DHCP Discover.
2. The BNG with the master subscriber routed redundancy protocol (SRRP) state or active MC-LAG link (called the session-owning node) will trigger the Diameter client to send a CCR-I.
3. The Diameter client generates a CCR-I and sends it to the active Diameter proxy, in this case using the internal Diameter connection.
4. The proxy sends the CCR-I to the PCRF.
5. The PCRF sends back the credit control answer CCA-I to the Diameter client via the Diameter proxy.
6. The Diameter proxy sends the CCA-I to the Diameter client using the internal Diameter connection.
7. The Diameter client sends a notification message to ESM.
8. ESM continues connecting the subscriber to the network.

Figure 94 **Subscriber Connecting to BNG Hosting Active Diameter Proxy**

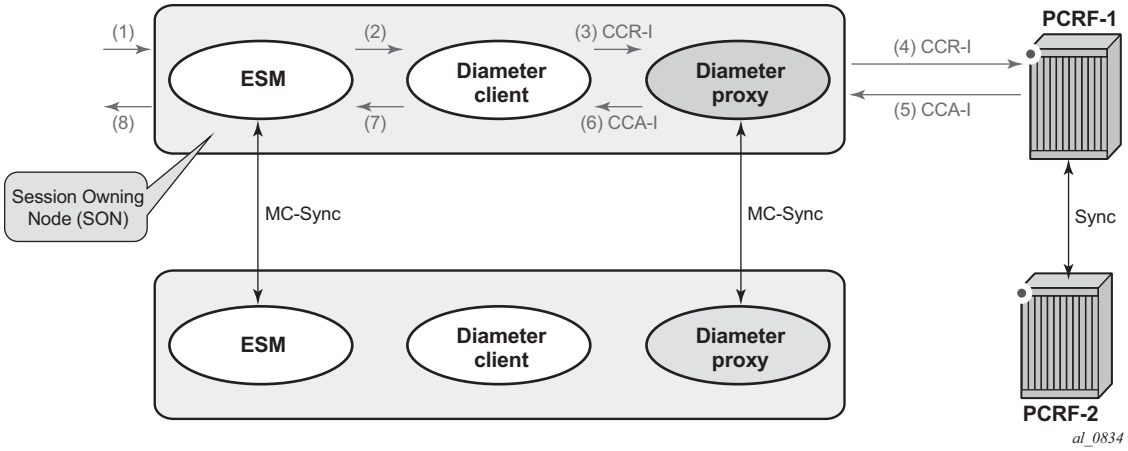

If the active proxy is not on the session-owning node, the scenario differs slightly, see Figure 95. In this case, the CCR-I in step 3 is sent from the session owning node to the proxy on the other BNG, and similarly, in Step 7, the CCA-I is sent from the proxy to the other BNG, using the external Diameter connection.

Figure 95 **Subscriber Connecting to BNG Hosting Standby Diameter Proxy**

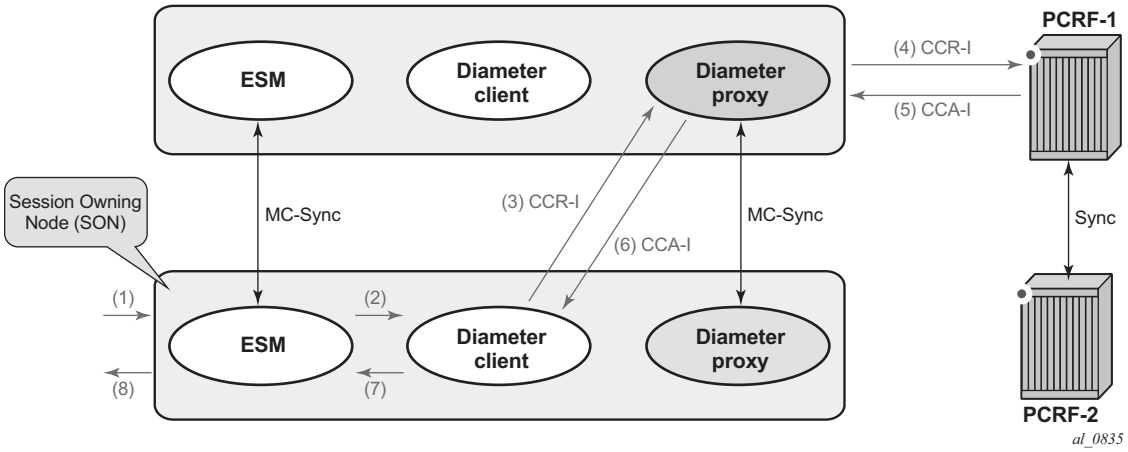

The Diameter proxy does not modify the Diameter packet except for the hop-by-hop identifier field in the Diameter header and the Origin-State-Id (OSI) AVP. This means that the Diameter client has to use the hostname/realm of the PCRF as the destination host/realm. Similarly, the PCRF will use the hostname/realm of the Diameter client as origin-host/realm, instead of the hostname/realm of the Diameter proxy. The exception will be the Diameter Capability messages, exchanged with the PCRF and originated by the Diameter proxy.

The Diameter proxy does not maintain any session state, so no session state information has to be synchronized with the other BNG. The Diameter proxy only synchronizes information that is needed for the Diameter connections to the network. This includes the chassis MAC address for active proxy selection and the OSI. Synchronizing the OSI results in a constant value for the OSI after a switchover to the other proxy; therefore, the Diameter server will not notice a restart of the BNG.

Although the Diameter proxy does not maintain any session state, it maintains the transaction state. This means that the transaction state is created when a request is sent to the network, and the state is removed when the answer is received and forwarded to the client, or when the transaction timer expires (by default 10 s) if no answer is received. This transaction state is maintained to know which client to send the answer messages to. For instance, when the Diameter proxy receives a CCA-I, it looks up the transaction state to find out which client originated the request. It then forwards the answer message to that client, and clears the transaction state. The handling of re-authorization request (RAR) messages is different because when a RAR message is received, there is no transaction state. In this case, the proxy will send the RAR message to both clients, and the client on the session-owning node will process this message and reply with a re-authorization answer (RAA) message. The client on the other BNG will ignore the RAR message.

IP Addressing

The Diameter client has one IP address, while the proxy has two IP addresses: one toward the client, and one toward the server. This is shown in Figure 96.

Figure 96 **Diameter IP Addressing**

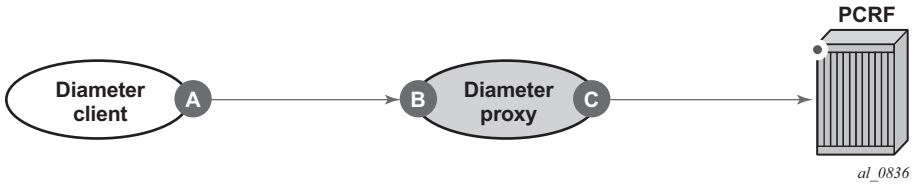

The Diameter client will use IP address A to originate Diameter connections toward the Diameter proxy, and will use IP address B as the IP destination address. The Diameter proxy will listen on IP address B for incoming Diameter connections, and will use IP address C as the IP source address for outgoing Diameter connections. All IP addresses can be in the same virtual private routed network (VPRN), or in a different VPRN, as long as IP routing from IP address A to IP address B is available. IP address B and C can also be the same IP address. The IP addresses must belong to an IP interface (for example, a loopback IP address).

Retransmissions

The Diameter proxy does not retransmit any messages with the exception of retransmissions at the TCP layer. Retransmissions at the Diameter level are done by the Diameter client; therefore, the proxy does not have to do these. When the client retransmits at the Diameter level, the T-bit (retransmission bit) is set in the Diameter header to indicate that it is a retransmitted message. The Diameter proxy will route messages based on the T-bit (see also Figure 97):

- If the proxy receives a Diameter message without the T-bit set, the message is forwarded to the primary Diameter connection.
- If the proxy receives a Diameter message with the T-bit set, the message is forwarded to the secondary Diameter connection.

Figure 97 **Retransmission Scenario**

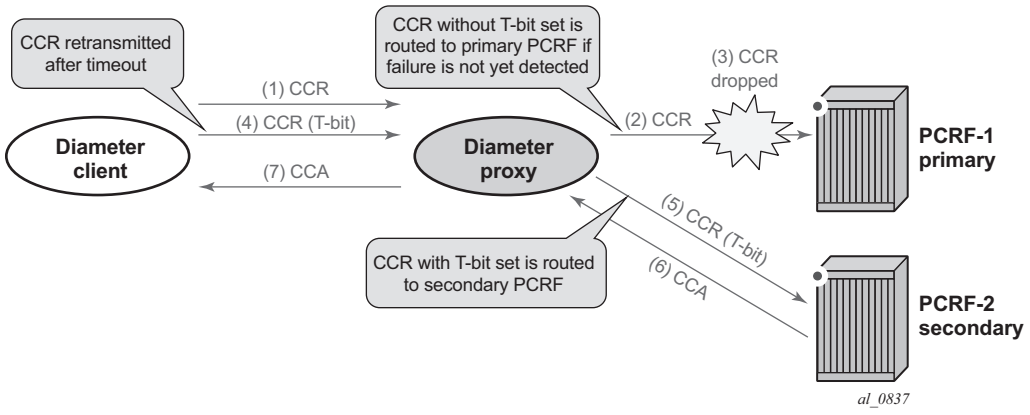

In the example of Figure 97, the client sends a CCR message to the proxy (1). The proxy forwards this CCR to the primary PCRF (2), but somewhere the message gets dropped (3). The client will timeout (based on the transaction timer) and retransmit the CCR with the T-bit set (4). The proxy forwards this CCR to the secondary PCRF (5), which will respond (6), and this response is forwarded to the client (7).

To generate a retransmitted message at the Diameter level on the Diameter client, the transaction timer has to expire and the failure handling parameter has to be set to *retry-and-terminate* or *continue*. For instance, in the case of Gx, both options will retransmit a CCR-I with the T-bit set. If no response to the retransmitted CCR-I message is received, failure handling *retry-and-terminate* will reject the host, while failure handling *continue* will accept the host. Note: The Diameter client retransmits

the same Diameter message to the same Diameter proxy only if there is just one Diameter peering connection open. This is typically the case in this scenario because the standby Diameter proxy rejects all incoming connections. If a secondary Diameter peering connection is open in the client, the retransmission is sent to that Diameter connection.

If the proxy receives a message with the T-bit set, and it only has a primary connection, then this message will be sent to the primary connection.

Proxy Switchover

A switchover to the other proxy is performed in the following cases:

- the BNG managing the active proxy goes down
- all Diameter connections on the active proxy go down

For fast convergence toward the new active proxy, it is advised to configure short-duration connection timers, transaction timers, and watchdog timers (see the Establishing a Diameter Peering Session chapter for an explanation of these timers). This will speed up the detection of failed Diameter connections, and also the establishment of new Diameter connections. The new proxy can only become active and fully operational when it has established at least one Diameter connection to the network, and both Diameter clients have detected the failed proxy and re-established Diameter connections to the new active proxy.

The scenario for a failing proxy is shown in Figure 98.

Figure 98 **Scenario with Failing Diameter Proxy**

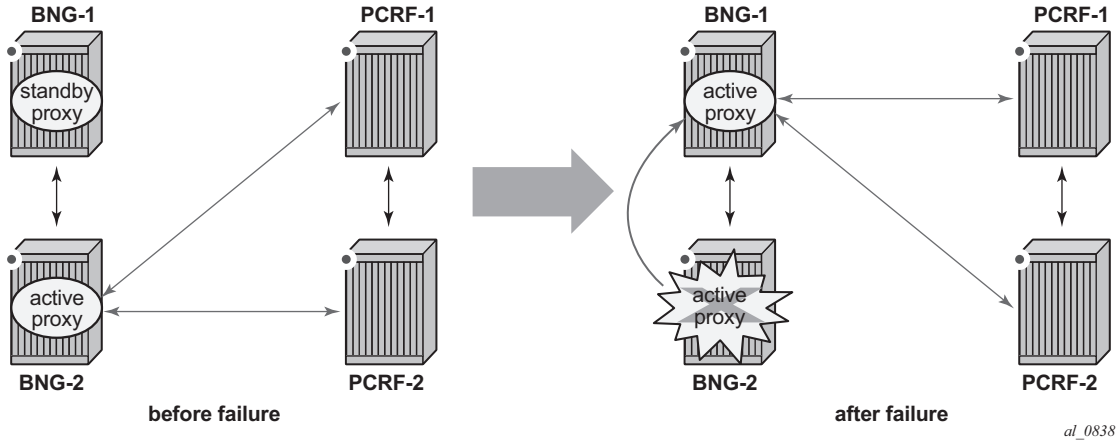

First, both BNGs are operational and BNG-2 is the active proxy while BNG-1 is the standby proxy. Only BNG-2 will establish Diameter connections to the servers (PCRFs in this example). When BNG-2 goes down, the PCRFs will detect the BNG failure (if watchdogs have been enabled in the PCRF) and remove the Diameter connection to the failed BNG. In parallel, BNG-1 detects that BNG-2 is down and starts establishing the Diameter connections to the PCRFs, and as soon as one of them is up, the BNG-1 proxy goes to active state.

Figure 99 and Figure 100 show the same scenario, but also for failing Diameter connections. Figure 99 shows that, after the failure of a Diameter connection to one PCRF, there is no switchover. When there is a second failure, while the Diameter connection to PCRF-1 is still down, a proxy switchover occurs although BNG-2 stays up, as shown in Figure 100.

Figure 99 **Scenario with One Diameter Connection Failing**

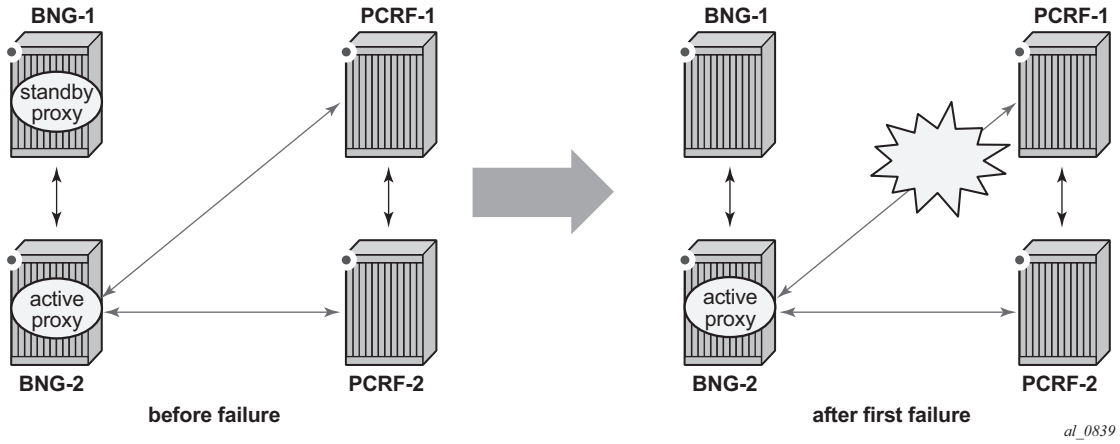

Figure 100 **Scenario with Both Diameter Connections Failing**

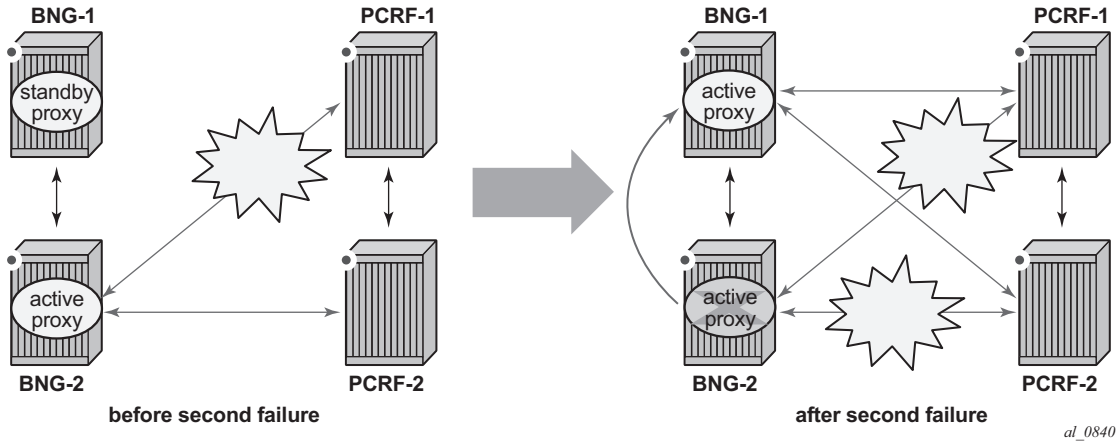

Access Node Gateway Change Trigger

In the case of Gx, the BNG originating the CCR-I message can optionally include the Access Node Gateway address (AN_GW_ADDRESS) containing the system IP address. Because there are now two BNGs, each having their own system IP address, the AN_GW_ADDRESS can change during the lifetime of a Diameter session. In Figure 101, BNG-1 is the ESM/SRRP master when the subscriber sends a DHCP discover. Consequently, BNG-1 originates the CCR-I with AN_GW_ADDRESS 10.0.0.2, which is its own system IP address. When BNG-1 fails, BNG-2 will take over and can optionally generate a CCR-U with the new AN_GW_ADDRESS, its own system IP address (10.0.0.11 in this example). BNG-2 will only do this if the PCRF has enabled event trigger 21 in the CCA-I. If the PCRF did not enable this event trigger, the PCRF will not be aware that BNG-2 became the new ESM/SRRP master. However, this should not have any effect on the correct behavior of the BNGs and the PCRF; therefore, this event trigger is an optional, informational event trigger. Note: Even without this event trigger, the PCRF will detect a failing Gx session to BNG-1 if watchdogs have been enabled in the PCRF.

Figure 101 **Access Node Gateway Change Trigger**

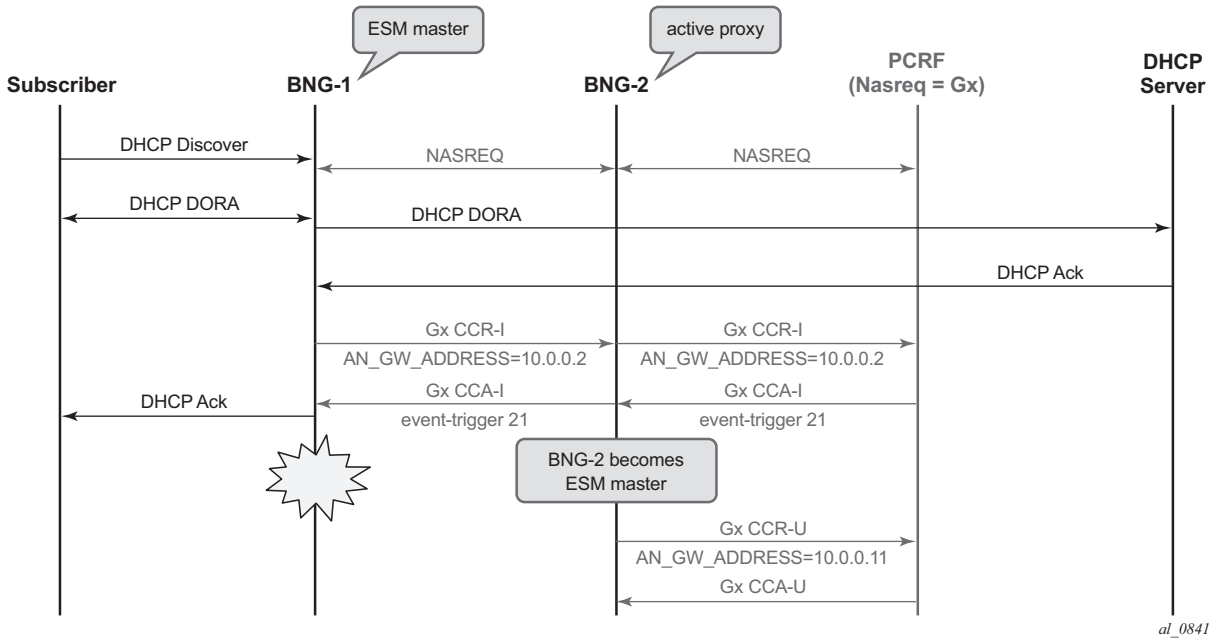

Configuration

To configure Diameter inter-chassis redundancy with Diameter proxy, four parts have to be configured:

1. multi-homed BNGs
2. Diameter client
3. Diameter proxy MC-Sync
4. Diameter proxy

The first two parts are common to ESM and to Diameter applications such as Gx and NASREQ, which means that these two first parts are not Diameter-proxy specific. The parts apply equally to single-chassis or dual-chassis configuration. The first part configures SRRP with MC-Sync and, optionally, a shunt tunnel.

The second part configures the Diameter client, which is unaware of whether it is interfacing directly with a Diameter proxy, an external DRA, or a PCRF. However, in multi-chassis redundancy with Diameter proxy, the destination host/realm and the IP addresses must be configured in the following way:

- The destination host/realm must be the AAA/PCRF host/realm and not the host/realm of the proxy.
- The destination IP address must be the IP address of the proxy and not the IP address of the AAA/PCRF.

An example configuration of two multi-homed BNGs follows. The realm of the AAA/PCRF is **dscrealm**. The IP address of the proxy on BNG-1 is 10.23.3.131 and the IP address of the proxy on BNG-2 is 10.23.4.131.

For example, the configuration of BNG-1 could be as follows:

BNG-1:

```
configure aaa diameter-peer-policy "DIAMETER.CLIENT"
   applications gx nasreq
   connection-timer 1
   origin-host "bngclient.geo.SRrealm"
   origin-realm "SRrealmClient"
   python-policy "dsc_subscriber_id"
   router 10000
   source-address 10.23.3.132
   peer "Proxy.Local" create
       address 10.23.3.131
       destination-realm "dscrealm"
       preference 10
       no shutdown
   exit
   peer "Proxy.Remote" create
```

```
        address 10.23.4.131
        destination-realm "dscrealm"
        preference 20
        no shutdown
    exit
```

For example, the configuration of the Diameter client on BNG-2 could be as follows:

BNG-2:

```
configure aaa diameter-peer-policy "DIAMETER.CLIENT"
    applications gx nasreq
    connection-timer 1
    origin-host "bngclient.geo.SRrealm"
    origin-realm "SRrealmClient"
    python-policy "dsc_subscriber_id"
    router 10000
    source-address 10.23.4.132
    peer "Proxy.Local" create
        address 10.23.4.131
        destination-realm "dscrealm"
        preference 10
        no shutdown
    exit
    peer "Proxy.Remote" create
        address 10.23.3.131
        destination-realm "dscrealm"
        preference 20
        no shutdown
    exit
```

Note: Both BNGs use the same origin-host (bngclient.geo.SRrealm) and the same origin-realm (**SRrealmClient**).

After the Diameter peer policy has been configured, it can be used in a Diameter application policy, which can be applied to the group interface or SAP. An example (minimal) configuration of a Gx Diameter application policy, using the previously defined Diameter peer policy with its assignment to a group interface, follows:

```
configure subscriber-mgmt diameter-application-policy "DIAMETER.Gx"
    description "Diameter application policy for Gx to Diameter Proxy"
    application gx
    diameter-peer-policy "DIAMETER.CLIENT"

configure service vprn 10001
    <…snip…>
    subscriber-interface "to_subscribers" create
      <…snip…>
      group-interface "to_STC" create
       <…snip…>
         diameter-application-policy "DIAMETER.Gx"
```

In the third configuration part, the MC-Sync for the Diameter proxy has to be enabled. This is done by adding **diameter-proxy** in the multi-chassis redundancy configuration, as shown in the following example for BNG-1 (10.0.0.2) and BNG-2 (10.0.0.11):

BNG-1:

```
configure redundancy
   multi-chassis
       peer 10.0.0.11 create
           source-address 10.0.0.2
           sync
               diameter-proxy
               srrp
               sub-mgmt ipoe pppoe
               port 3/2/4 sync-tag "stc" create
               exit
               no shutdown
           exit
           no shutdown
       exit
   exit
```

BNG-2:

```
configure redundancy
   multi-chassis
       peer 10.0.0.2 create
           source-address 10.0.0.11
           sync
               diameter-proxy
               srrp
               sub-mgmt ipoe pppoe
               port 1/1/1 sync-tag "stc" create
               exit
               no shutdown
           exit
           no shutdown
       exit
   exit
```

In the preceding configuration, only the **diameter-proxy** command is specific to the Diameter proxy configuration and all the rest is part of the regular multi-homed BNG configuration in part 1.

The fourth part is the actual Diameter proxy configuration and it re-uses the Diameter client configuration commands with three exceptions:

- when creating the Diameter peer policy, the role “proxy” has to be enabled.
- the Diameter peer policy of the proxy does not have to be assigned to any group interface or SAP.

- proxy configuration has to be added, which contains the MCS-Peer IP address referring to the configuration in the previous part, and the IP address and VPRN for incoming Diameter connections.

For example, the Diameter proxy configuration for BNG-1 and BNG-2 could be as follows:

BNG-1:

```
configure aaa diameter-peer-policy "DIAMETER.PROXY" role proxy create
   applications gx nasreq
   connection-timer 1
   origin-host "bng.geo.SRrealm"
   origin-realm "SRrealm"
   router 10000
   source-address 10.23.3.130
   peer "DSC.39.118" create
       address 10.40.15.5
       destination-realm "dscrealm"
       preference 20
       no shutdown
   exit
   peer "DSC.39.122" create
       address 10.40.15.4
       destination-realm "dscrealm"
       preference 10
       no shutdown
   exit
   proxy
       mcs-peer  10.0.0.11 sync-tag "stc"
       router 10000 address 10.23.3.131
       no shutdown
   exit
```

BNG-2:

```
configure aaa diameter-peer-policy "DIAMETER.PROXY" role proxy create
   applications gx nasreq
   connection-timer 1
   origin-host "bng.geo.SRrealm"
   origin-realm "SRrealm"
   router 10000
   source-address 10.23.4.130
   peer "DSC.39.118" create
       address 10.40.15.5
       destination-realm "dscrealm"
       preference 20
       no shutdown
   exit
   peer "DSC.39.122" create
       address 10.40.15.4
       destination-realm "dscrealm"
       preference 10
       no shutdown
   exit
   proxy
```

```
        mcs-peer  10.0.0.2 sync-tag "stc"
        router 10000 address 10.23.4.131
        no shutdown
    exit
```

Optionally, Python can be applied to all Diameter connections (see Figure 102):

- at the Diameter client side on the Diameter connection to the proxy.
- at the Diameter proxy side on the Diameter connection coming from the client.
- at the Diameter proxy side on the Diameter connection going to the network.

Figure 102 **Python for Diameter Connections**

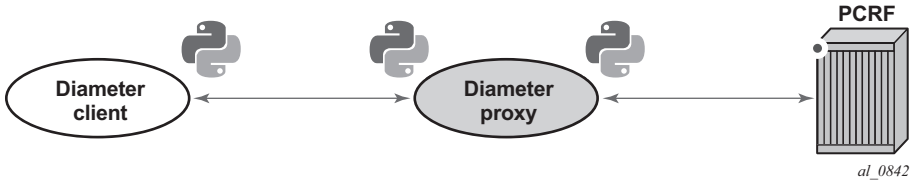

When using Python, the script has to be applied to the correct message type and direction. For example, if the following Python policy is applied to the proxy Diameter peer policy, the scripts are applied to the CCR messages toward the PCRF (egress direction) and to the CCA messages toward the client (egress direction), see Figure 103:

```
python-policy "change_ipcan" create
   diameter ccr direction egress script "change_ipcan"
   diameter cca direction egress script "change_ipcan_answer"
exit
```

Figure 103 **Python Example**

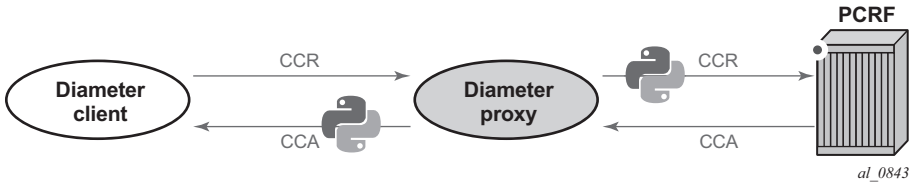

Troubleshooting

The state of the Diameter client can be verified with the following command:

```
# show aaa diameter-peer-policy "DIAMETER.CLIENT"
===============================================================================
```

```
Diameter Peer Policy : DIAMETER.CLIENT
===============================================================================
Last Mgmt Change      : 06/16/2015 13:05:25
Applications          : gx nasreq
Role                  : client
Description           : (Not Specified)
-------------------------------------------------------------------------------
Diameter Config Values
-------------------------------------------------------------------------------
Origin Host           : bngclient.geo.SRrealm
Origin Realm          : SRrealmClient
Connection Timer      : 1                   Source Address     : 10.23.3.132
Transaction Timer     : 30 (default)        Router             : 10000
Watchdog Timer        : 30 (default)
Vendor Support        : 3GPP (default)
Python Policy         : dsc_subscriber_id
-------------------------------------------------------------------------------
Peer Name        Oper  PSM State     Susp  Cooldown  Pref  Order  Pri/Sec
-------------------------------------------------------------------------------
Proxy.Local      Yes   Wait-Conn-Ack No    -         10    -      -
Proxy.Remote     Yes   I-Open        No    -         20    1      Primary
===============================================================================
```

Because the Diameter client is unaware that it is interfacing with a proxy, there is no proxy-specific information in this output. However, the standby proxy does not accept incoming Diameter connections; therefore, the connection on the client to this standby proxy will not be in I-Open state. In the preceding example, the peer “Proxy.Remote” is in I-Open state, indicating that this is the peer toward the active proxy (which is on BNG-2 in this example).

The CLI command was taken on BNG-1. Similar information can be seen on BNG-2, but there the Diameter connection toward the local proxy is in I-Open state:

```
# show aaa diameter-peer-policy "DIAMETER.CLIENT"
===============================================================================
Diameter Peer Policy : DIAMETER.CLIENT
===============================================================================
Last Mgmt Change      : 06/16/2015 12:47:14
Applications          : gx nasreq
Role                  : client
Description           : (Not Specified)
-------------------------------------------------------------------------------
Diameter Config Values
-------------------------------------------------------------------------------
Origin Host           : bngclient.geo.SRrealm
Origin Realm          : SRrealmClient
Connection Timer      : 1                   Source Address     : 10.23.4.132
Transaction Timer     : 30 (default)        Router             : 10000
Watchdog Timer        : 30 (default)
Vendor Support        : 3GPP (default)
Python Policy         : dsc_subscriber_id
-------------------------------------------------------------------------------
Peer Name        Oper  PSM State     Susp  Cooldown  Pref  Order Pri/Sec
-------------------------------------------------------------------------------
Proxy.Local      Yes   I-Open        No    -         10    1      Primary
Proxy.Remote     Yes   Wait-Conn-Ack No    -         20    -      -
```

```
===============================================================================
```

The Diameter proxy can be verified with a similar CLI command (see the following). There are three parts in this CLI command that are specific for the proxy:

- Diameter proxy shows the administrative state and operational state. The out-of-service state means that the proxy is not active, while the in-service state means that the proxy is operational and active, that is, establishing Diameter connections and processing Diameter messages. The out-of-service state means that the proxy is either not operational or that it is standby until the other proxy fails. For the rest, this part shows the MCS specific information for the proxy.
- Proxy multi-chassis redundancy shows which proxy is the active one (the local or the remote proxy). The possible states of the proxy are:
  - Active — The proxy is active; this also means that it has at least one Diameter connection in I-Open state.
  - activeWait — The proxy is selected to become active, but is still setting up the Diameter connections with the Diameter server(s).
  - standbyWait — The proxy is waiting for its multi-chassis synchronization (MCS) partner to become active; a proxy cannot become standby as long the other side is not active.
  - proxySwitchoverReq — The proxy was active and is becoming standby.

Besides the state information, this part also shows the synchronized information (MAC address and OSI). The controller is the BNG with the highest MAC address (independent of whether it is active or standby), and the controller MAC address is its MAC address. Because MAC addresses do not change, the controller also does not change, except if the controller goes down, or a BNG detects a higher MAC address than the current one. If a BNG does not detect another BNG, that BNG declares itself as controller. Because the controller has the highest MAC address, it is the controller that becomes active when there is no active proxy or there is a conflict and a new selection has to be made.

- Client-side peers show the incoming Diameter connections from the clients. On the BNG with the standby proxy, there will be no incoming Diameter connections (in this example, BNG-1 has the standby proxy, so no client-side peers).

BNG-1:

```
# show aaa diameter-peer-policy "DIAMETER.PROXY"
===============================================================================
Diameter Peer Policy : DIAMETER.PROXY
===============================================================================
Last Mgmt Change      : 06/16/2015 13:05:25
Applications          : gx nasreq
Role                  : proxy
Description           : (Not Specified)
```

```
-------------------------------------------------------------------------------
Diameter Config Values
-------------------------------------------------------------------------------
Origin Host           : bng.geo.SRrealm
Origin Realm          : SRrealm
Connection Timer      : 1                 Source Address     : 10.23.3.130
Transaction Timer     : 30 (default)      Router             : 10000
Watchdog Timer        : 30 (default)
Vendor Support        : 3GPP (default)
Python Policy         : N/A
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Diameter proxy
-------------------------------------------------------------------------------
Proxy administrative state  : enabled
Proxy operational state     : out-of-service
Router                      : 10000
IP address                  : 10.23.3.131
MCS peer IP address         : 10.0.0.11
MCS sync tag                : stc
Last management change      : 06/16/2015 13:05:25
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Proxy multi-chassis redundancy
-------------------------------------------------------------------------------
                                  Local               Remote
-------------------------------------------------------------------------------
State                             standby             active
Origin-State-Id                   1434460525          1434460525
MAC address                       00:03:fa:14:2b:a7   00:21:05:9b:b8:24
Controller MAC address            00:21:05:9b:b8:24   00:21:05:9b:b8:24
Controller                        No                  Yes
-------------------------------------------------------------------------------
Peer Name            Oper  PSM State     Susp  Cooldown  Pref  Order  Pri/Sec
-------------------------------------------------------------------------------
DSC.39.118           Yes   Closed        No    -         20    -      -
DSC.39.122           Yes   Closed        No    -         10    -      -
No client-side peers found.
===============================================================================
```

BNG-2:

```
# show aaa diameter-peer-policy "DIAMETER.PROXY"
===============================================================================
Diameter Peer Policy : DIAMETER.PROXY
===============================================================================
Last Mgmt Change      : 06/16/2015 12:47:14
Applications          : gx nasreq
Role                  : proxy
Description           : (Not Specified)
-------------------------------------------------------------------------------
Diameter Config Values
-------------------------------------------------------------------------------
Origin Host           : bng.geo.SRrealm
Origin Realm          : SRrealm
Connection Timer      : 1                 Source Address     : 10.23.4.130
Transaction Timer     : 30 (default)      Router             : 10000
Watchdog Timer        : 30 (default)
```

```
Vendor Support        : 3GPP (default)
Python Policy         : N/A
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Diameter proxy
-------------------------------------------------------------------------------
Proxy administrative state  : enabled
Proxy operational state     : in-service
Router                      : 10000
IP address                  : 10.23.4.131
MCS peer IP address         : 10.0.0.2
MCS sync tag                : stc
Last management change      : 06/16/2015 12:47:14
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Proxy multi-chassis redundancy
-------------------------------------------------------------------------------
                                    Local                Remote
-------------------------------------------------------------------------------
State                               active               standby
Origin-State-Id                     1434460525           1434460525
MAC address                         00:21:05:9b:b8:24    00:03:fa:14:2b:a7
Controller MAC address              00:21:05:9b:b8:24    00:21:05:9b:b8:24
Controller                          Yes                  No
-------------------------------------------------------------------------------
Peer Name            Oper  PSM State     Susp  Cooldown  Pref  Order  Pri/Sec
-------------------------------------------------------------------------------
DSC.39.118           Yes   I-Open        No    -         20    2      Secondary
DSC.39.122           Yes   I-Open        No    -         10    1      Primary
-------------------------------------------------------------------------------
Client-side peers
-------------------------------------------------------------------------------
IP address      TCP port   PSM state
-------------------------------------------------------------------------------
10.23.3.132     58664      r-open
10.23.4.132     61960      r-open
===============================================================================
```

The preceding command shows the IP address and TCP port of the incoming Diameter connections. The following command can be used to get the statistics of these connections:

```
# show aaa diameter-peer-policy "DIAMETER.PROXY" statistics client-side-peer-
ip 10.23.3.132 port 58664
===============================================================================
Client-side peer statistics
===============================================================================
IP address               : 10.23.3.132
TCP port                 : 58664
time statistics cleared  : 06/17/2015 12:21:36
===============================================================================
Client initiated tx/rx                  Server initiated tx/rx
-------------------------------------------------------------------------------
TCP Send Failed          : 0            TCP Send Failed          : 0
Diam Rx Drop Count (Reqs) : 0           Diam Rx Drop Count (Resps): 0
Diam Rx Requests         : 13818        Diam Tx Requests         : 21085
Diam Tx Responses        : 13818        Diam Rx Responses        : 21085
Pending Messages         : 0
```

```
Lifetime Timeouts         : 0

Client initiated tx/rx                  Server initiated tx/rx
-------------------------------------------------------------------------------
CCR initial Rx            : 1           CCA initial Tx            : 1
CCR update Rx             : 1           CCA update Tx             : 1
CCR terminate Rx          : 1           CCA terminate Tx          : 1
CER Rx                    : 1           CEA Tx                    : 1
DWR Rx                    : 13814       DWA Tx                    : 13814
DWR Tx                    : 21085       DWA Rx                    : 21085
ASR Tx                    : 0           ASA Rx                    : 0
RAR Tx                    : 0           RAA Rx                    : 0
DPR Tx                    : 0           DPA Rx                    : 0
DPR Rx                    : 0           DPA Tx                    : 0
AAR Rx                    : 0           AAA Tx                    : 0
```

The output shows the number of messages per message type that have been sent or received.

Because the Diameter proxy uses the chassis MAC address, it might be useful to view this MAC address. This address can be shown with the following command:

BNG-1:

```
# show chassis
   --- snipped ---
   Base MAC address : 00:03:fa:14:2b:a7
```

BNG-2:

```
# show chassis
   --- snipped ---
   Base MAC address : 00:21:05:9b:b8:24
```

Debug is available at all of the following connections (see Figure 104):

- at the Diameter client side on the Diameter connection to the proxy.
- at the Diameter proxy side on the Diameter connection coming from the client.
- at the Diameter proxy side on the Diameter connection going to the network.

Figure 104 **Diameter Debugging Points**

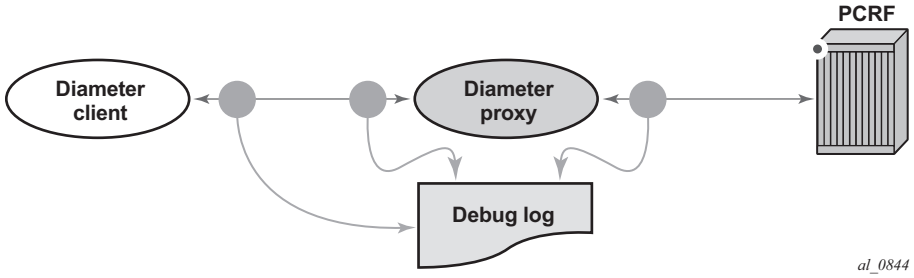

These three connections mean that, if all debugs are enabled, all messages are displayed three times in the debug log. The same debug commands as for the client apply to the proxy:

```
# debug diameter
  - diameter
  - no diameter

 [no] dest-realm      - Restrict output to a specific destination-realm
      detail-level    - Configure the detail level of debug output
 [no] diameter-peer   - Restrict output to a specific peer
 [no] diameter-peer-* - Restrict output to a specific policy
 [no] message-type    - Restrict output to specific message-types
 [no] origin-realm    - Restrict output to a specific origin-realm
```

The following traps are generated by the Diameter proxy:

- tmnxDiamPpPrxMcLocStateChanged is generated each time the proxy changes state.
- tmnxDiamPrxMessageDropped is generated each time the proxy drops a message because it is unable to process it.

```
# show log event-control "diameter"
Application
 ID#    Event Name                          P   g/s     Logged     Dropped
-------------------------------------------------------------------------------
   2001 tmnxDiamPolicyPeerStateChange    MI  thr      38443           0
   2002 tmnxDiamAppMessageDropped        MI  thr          0           0
   2003 tmnxDiamAppSessionFailure        MI  thr          0           0
   2004 tmnxDiamSessionEvent             MI  thr          0           0
   2005 tmnxDiamPpPrxMcLocStateChanged   MI  thr          3           0
   2006 tmnxDiamPrxMessageDropped        MI  thr          0           0
===============================================================================
```

Conclusion

This chapter describes the Diameter proxy, which can be used in case a multi-homed BNG configuration (inter-chassis redundancy) is used. The proxy ensures that both BNGs use the same Diameter hostname/realm and that the AAA/PCRF in the network only sees one logical BNG. The benefit is that all messages belonging to the same Diameter session will have the same origin-host/realm, irrespective of which BNG is active.

ESM Basics

This chapter provides information about Enhanced Subscriber Management.

Topics in this chapter include:

Applicability

This chapter is applicable to 7450 ESS, 7750 SR and 7710 SR series and was tested on SR OS 11.0.R4. Chassis mode b or higher must be used.

Summary

Subscriber management in general includes the following functions:

- subscriber host authentication, identification, addressing, authorization and accounting:
  - authentication – Check whether the subscriber host is allowed access via a Local User Data Base (LUDB) or via a RADIUS server.
  - identification – Fetch the data to use for the subscriber host, including the definition of the subscriber-ID.
  - addressing – Fetch the address information to use, IPv4 and/or IPv6.
  - authorization – Check what the subscriber host is allowed to do.
  - accounting – Both off-line charging (RADIUS and XML) and on-line charging (RADIUS credit control and Diameter credit control application) are supported.
- subscriber host instantiation, based on:
  - A protocol (DHCP, PPPoE/oA/oEoA or ARP) for dynamic hosts, and started through a trigger packet.
  - A static configuration for static hosts.

- subscriber QoS - Ensure per service and per application SLAs, based on:
  - the overall subscriber rate
  - subscriber profiles
  - service level agreement profiles
- subscriber security:
  - Avoid malicious access through anti-spoofing and based on access control lists (ACLs) / IP-filters.
  - DDOS mitigation.
- subscriber persistency – The subscriber state is written to flash-disk (a.k.a. persistent data) and automatically restored on node reboot.
- subscriber resiliency – The subscriber state is retained in case of a node failure in redundant environments, through Multi Chassis Synchronization (MCS).
- subscriber troubleshooting, using:
  - OAM test.
  - mirroring.
  - debugging and event logging.

Enhanced Subscriber Management (ESM) implies subscriber management functions are applied at subscriber level. Subscriber hosts are created, and all of the features listed above apply.

Basic Subscriber Management (BSM) implies subscriber management functions are applied at SAP level. Only a subset of the functions listed above apply.

This example gives an overview of the ESM data required to perform subscriber management functions, and how the ESM data is organized.

Overview

The following terminology is used extensively in Triple Play Service Delivery Architecture (TPSDA) and is key to understanding ESM:

- device
- subscriber host
- subscriber

A device is equipment located at the customer premises. Example devices are computers, smart-phones, set-top boxes, etc including the Residential Gateway (RGW) providing the connection towards the Internet. The RGW is connected to the access network using an XDSL-connection, a PON-connection, etc.

A subscriber is a collection of subscriber hosts connected to a single RGW. The subscriber is identified by its subscriber-ID, a character string of 16 characters maximum which is used for administrative purposes.

Unlike devices, which are physical entities, subscribers and subscriber hosts are logical entities. These are created dynamically and resources are allocated when a device connects to the network and becomes active.

The following host types are recognized by the Broadband Network Gateway (BNG):

- DHCP hosts
- PPPoE hosts
- ARP hosts

The BNG uses the combination of following parameters to uniquely identify a single subscriber host:

- SAP-id
- IP-address
- MAC-address
- session-ID (PPPoE only)

Multiple subscriber hosts can be associated with a single device.

Figure 105 shows a bridged RGW scenario. Subscriber-1has two devices, device-1 and device-2. Device-1 is a dual stack PPPoE device and is assigned an IPv4 address for host-1 and an IPv6 SLAAC prefix for host-2, both running over a single PPPoE session. Device-2 is a single stack IPoE device and is assigned an IPv4 address for host-3.

Figure 105 **Bridged RGW Scenario**

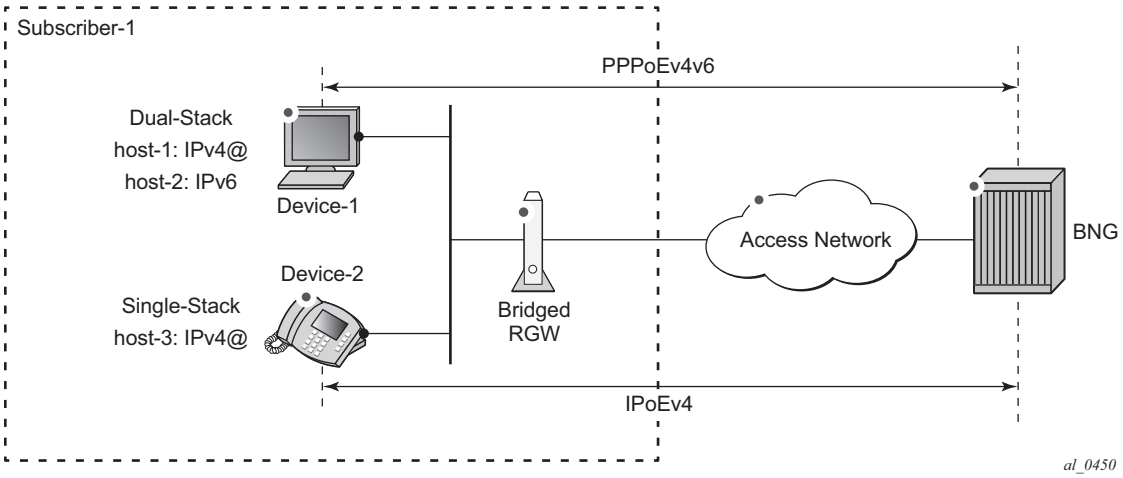

Figure 106 shows a routed dual stack RGW scenario. Subscriber-2 has two devices but they are hidden behind the RGW. The IP-addresses used by these devices are not known by the BNG. The RGW contains three hosts:

- Host-1 is assigned an outside IPv4 address, which the RGW uses for Network Address Translation (NAT).
- Host-2 is assigned an IPv6 wan host address or SLAAC prefix, which the RGW uses towards the outside network.
- Host-3 is assigned an IPv6 prefix for prefix delegation, which the RGW uses for allocating IPv6 addresses to IPv6 capable devices in the home network.

Figure 106 **Routed RGW Scenario**

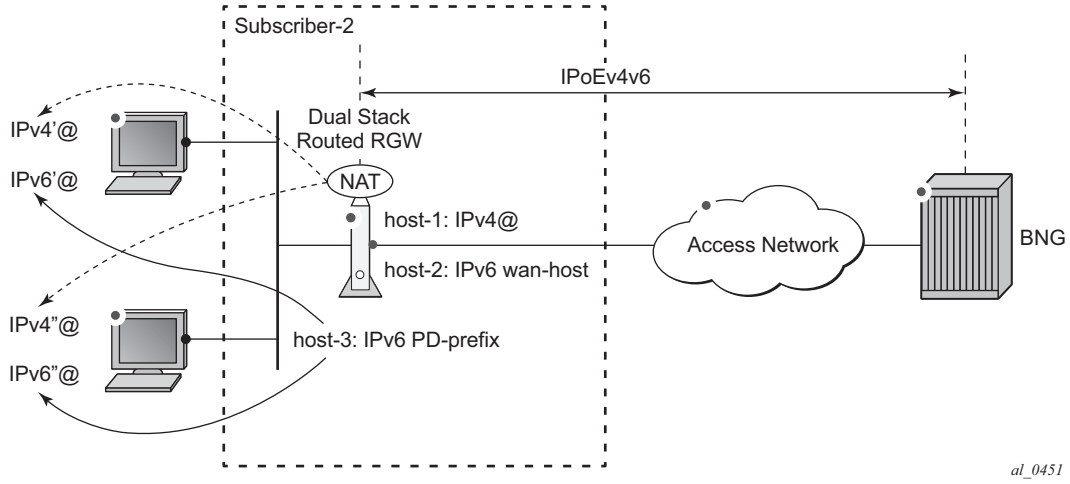

Detailed information on both scenarios can be found in ESMv6: IPoE Dual Stack Hosts and the ESMv6: PPPoE Dual Stack Hosts respectively.

Configuration

Figure 107 shows two objects closely related to subscribers and subscriber hosts:

- SLA profile
- Subscriber profile

Figure 107 **SLA-Profile and Sub-Profile**

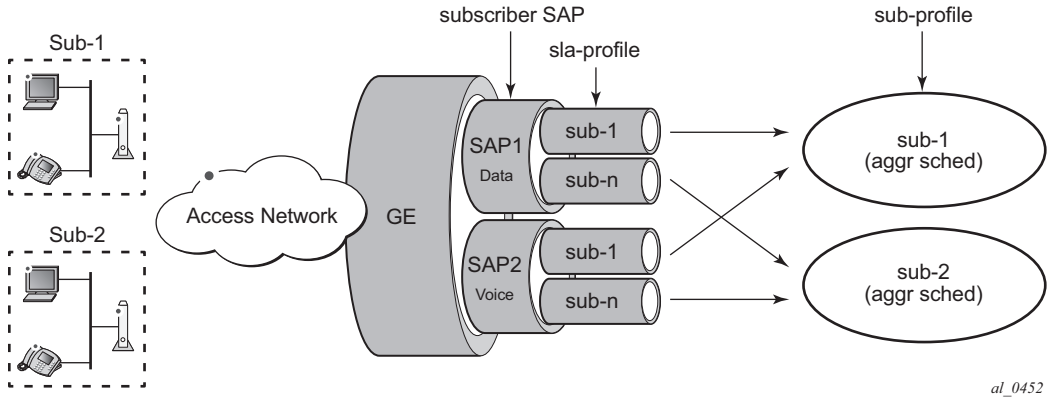

SLA Profile

A Service Level Agreement profile (SLA profile) is a template identified by name (maximum 32 chars) which defines:

- per service QoS settings as part of the QoS policy (ingress and egress):
  - queue/policer
  - scheduling/priority levels
  - bandwidth limits (CIR/PIR) and queue/policer depths (MBS/CBS)
  - classification
  - (re-)marking
- IP-filters for IPv4 and/or IPv6 (ingress and egress):
  - Access Control Lists (ACL)/Filters
- host-limit

- credit-control

An example SLA profile is shown below.

```
configure
   subscriber-mgmt
       sla-profile "sla-prof-1" create
           ingress
               qos 100
               exit
           exit
           egress
               qos 100
               exit
           exit
        exit
```

An instance of the sla-profile is created at host instantiation time.

The SLA profile enforces traffic control on a per service per subscriber basis. When multiple hosts for the same subscriber use the same SLA profile they share the same set of queues/policers as long as they are on the same SAP. The combined traffic for these hosts is controlled by the settings defined in the SLA profile.

Subscriber Profile

The subscriber profile (sub-profile) is a template identified by name (maximum 32 chars) which defines:

- per subscriber QoS settings:
  - aggregate rate (egress only)
  - scheduler policy (ingress and egress)
  - policer control policy (ingress and egress)
- the accounting profile (RADIUS or XML)
- multicast parameters (igmp-policy, etc.)
- Network Address Translation (NAT) parameters
- the ANCP (Access Node Control Protocol) parameters
- the default sla-profile mappings
- etc.

An example subscriber profile is shown below.

```
configure
    subscriber-mgmt
        sub-profile "sub-prof-1" create
```

```
            ingress
                policer-control-policy "pol-ctrl-1"
                exit
            exit
            egress
                scheduler-policy "down-1"
                exit
            exit
        exit
```

An instance of the sub-profile is created at host instantiation time.

The sub-profile enforces traffic control on a per subscriber basis. he aggregate traffic of all hosts of a particular subscriber is controlled by the settings defined in the sub-profile, as long as the subscriber hosts are terminated on the same card.

QoS details are out of the scope of this example.

Subscriber SAP

A subscriber SAP is a SAP in a VPLS (Bridged CO model, which is not covered), VPRN or IES service (both in the Routed CO model) on which the queues/policers and other resources are allocated and de-allocated on a per subscriber basis.

A static SAP is a subscriber SAP when sub-sla-mgmt is enabled and becomes operational when in the no shutdown state. Multiple subscribers can connect through this SAP simultaneously when multi-sub-sap is enabled.

```
configure
    service
        ies 1
            subscriber-interface "sub-int-1"
                group-interface "grp-int-1-1"
                    sap 1/1/1:101 create
                        description "data SAP for DLSAM-1"
                        sub-sla-mgmt
                            multi-sub-sap 400
                            --- snip ---
                            no shutdown
                        exit
                    exit
```

Single SAP parameters (profiled-traffic-only and non-sub-traffic) are taken into account only when multi-sub-sap is left at its default of 1.

Subscriber Host Identification and Instantiation

The general subscriber host identification (authentication) process is depicted in Figure 108 and encompasses fetching the ESM data for a connecting device.

Figure 108 **Subscriber Host Identification and Instantiation Process**

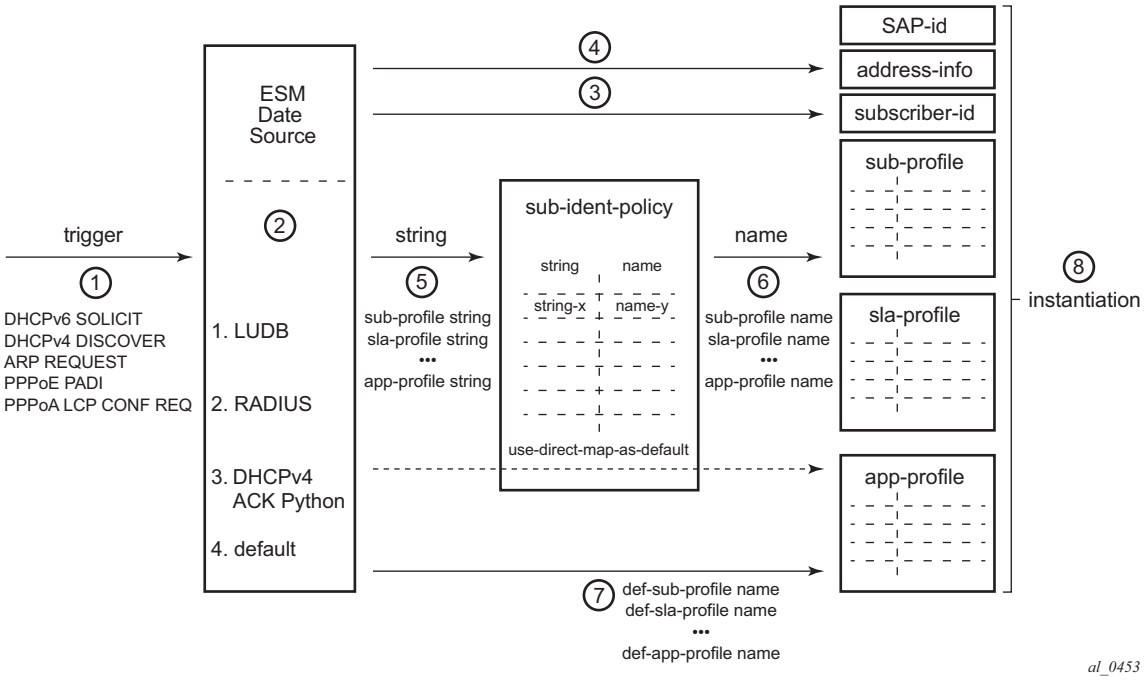

ESM data is organized as a set of profiles which control the behavior of individual subscribers and subscriber hosts.

Static hosts are instantiated using the manually provisioned ESM data.

For dynamic hosts, the host identification and instantiation process is started by a trigger message (1 in Figure 108) which can be one of the following:

- DHCPv4 DISCOVER
- DHPCv6 SOLICIT
- PPPoE PADI
- PPPoA LCP Configuration Request
- ARP REQUEST

The following types of ESM data are distinguished as mandatory or optional ESM data.

Mandatory ESM data is:

- address information (IPv4 and/or IPv6)
- subscriber ID
- sub-profile data
- sla-profile data

Optional ESM data includes:

- app-profile data
- inter-dest-id
- Access Node Control Protocol (ANCP) data

A host is instantiated only when all mandatory ESM data is available for that host.

The ESM data sources (2) are consulted in following predefined sequence:

1. LUDB
2. RADIUS
3. DHCPv4 ACK python
4. Default

The ESM data sources provide:

- the subscriber-id (3)
- the profile strings (5)

Address information (IPv4 and/or IPv6) (4) can be provided by:

- the LUDB or RADIUS ESM data sources
- a DHCP server

The LUDB, RADIUS or DHCPv4 ACK python ESM data sources provide profile strings (5) which have a maximum length of 16 characters. They must be translated into profile names (6) which have a maximum length of 32 characters. The profile name is the key to access the actual profile data. The ESM data source *default* directly returns profile names (7), which do not need any translation.

The ESM string to profile name translation is defined in configurable mapping tables which are part of the subscriber identification policy (sub-ident-policy). Mapping tables can be defined for:

- the sub-profile
- the sla-profile

- the app-profile

In case no mapping is needed because the strings and names are set to the same set of values, then a subscriber identification policy is needed with the attribute *use-direct-map-as-default*. (See section Subscriber Identification Policy.)

Note the subscriber-id does not need to be translated.

The instantiation process (8) ensures the subscriber host is created:

- The subscriber host is added to the active subscriber list.
- Resources (queues, policers, filters, etc) are allocated for SLA enforcement.
- Status information is updated.

As a result the system starts forwarding user data to and from that subscriber host.

A subscriber is instantiated as soon as the first subscriber host for this subscriber is instantiated, and deleted when the last subscriber host for this subscriber is deleted.

Subscriber Identification Policy

The subscriber identification policy is identified by name and defines:

- a description (optional)
- the sub-profile map
- the sla-profile map
- the app-profile map (optional)
- the location of a DHCPv4 ACK python script (optional):
  - primary, secondary and tertiary locations are possible
- the DHCP option used to get the identification strings from (optional)

A subscriber identification policy is needed for dynamic hosts only. In that case one of the ESM data sources from Figure 108 is used.

Using a python script for subscriber host identification and instantiation is restricted to IPv4 hosts when triggered by the DHCPv4 ACK message.

The first example has no explicit mappings.

```
configure
    subscriber-mgmt
        sub-ident-policy "sub-id-pol-1" create
            description "direct mapping policy"
```

```
            sub-profile-map
                use-direct-map-as-default
            exit
            sla-profile-map
                use-direct-map-as-default
            exit
            app-profile-map
                use-direct-map-as-default
            exit
        exit
```

The second example contains explicit mappings.

```
configure
    subscriber-mgmt
        sub-ident-policy "sub-id-pol-2" create
            description "explicit mapping policy"
            sub-profile-map
                entry key "sub-string-1" sub-profile "sub-prof-1"
                entry key "sub-string-2" sub-profile "sub-prof-2"
                entry key "sub-string-3" sub-profile "sub-prof-3"
            exit
            sla-profile-map
                entry key "sla-string-1" sla-profile "sla-prof-1"
                entry key "sla-string-2" sla-profile "sla-prof-2"
                entry key "sla-string-21" sla-profile "sla-prof-20"
                entry key "sla-string-22" sla-profile "sla-prof-20"
            exit
        exit
```

Note that multiple strings can be mapped to the same profile, as the example above shows.

The subscriber identification policy is applied at SAP level in the sub-sla-mgmt context, as shown below.

```
configure
    service
        ies 1
            subscriber-interface "sub-int-1" create
                --- snip ---
                group-interface "grp-int-1-1" create
                    --- snip ---
                    sap 1/1/1:111 create
                        sub-sla-mgmt
                            def-sub-profile "sub-prof-1"
                            def-sla-profile "sla-prof-1"
                            sub-ident-policy "sub-id-pol-1"
                            multi-sub-sap 400
                            no shutdown
                        exit
                    exit
                exit
            exit
```

Combining Multiple ESM Data Sources

Multiple ESM data sources from Figure 108 can be consulted for instantiating a subscriber host. When multiple ESM data sources are used, they are accessed in following sequence:

1. LUDB
2. RADIUS
3. DHCPv4 ACK python
4. Default

The outputs from the different data sources are merged. Data is appended as explained in the Flexible Authentication Model in ESM section of this guide.

Default ESM Data Profiles

Default ESM data profiles are used when the other ESM data sources (LUDB, RADIUS, DHCPv4 ACK python) only provide partial data, or no data at all. The default data complements the data provided by the data sources in order to get the full set of ESM data needed for host instantiation.

Default ESM data profiles are defined per SAP in the sub-sla-mgmt context, as shown in the following example.

```
configure
    service
        ies 1
            subscriber-interface "sub-int-1-1"
                group-interface "grp-int-1-1"
                    sap 1/1/1:111 create
                        sub-sla-mgmt
                            def-sub-profile "sub-prof-1"
                            def-sla-profile "sla-prof-1"
                            sub-ident-policy "sub-id-pol-1"
                            multi-sub-sap 400
                            no shutdown
                        exit
```

Address Information

Address information for both IPv4 as well as IPv6 is provided to subscriber hosts using a DHCPv4/v6 server, an LUDB or a RADIUS server.

The IPv4 and IPv6 address information encompasses:

- IPv4 and/or IPv6 address/mask and prefix
- default gateway (in IPv6 announced through RA messages)
- DNS server(s)
- etc

The following cases can be distinguished to obtain an IPv4 or IPv6 address/prefix upon LUDB/RADIUS authentication:

- LUDB/RADIUS returns a unique IPv4/IPv6 address/prefix, directly.
- LUDB/RADIUS returns a pool name, which then is resolved to a unique IP address/prefix through:
  - a DHCP server,
  - or Local Address Assignment (LAA).
- LUDB/RADIUS does not return a unique IPv4/IPv6 address/prefix, nor a pool-name. In that case a DHCP server is contacted using a gi-address or link-address to obtain an address/prefix.

Configuring a DHCP server is out of the scope of this example.

Direct Address Assignment using LUDB/RADIUS

In the example shown in Figure 109, a unique IP address is retrieved from LUDB/RADIUS. No interaction with a DHCP server is needed at all.

This scenario applies to IPoE (DHCPv4, DHCPv6 and SLAAC) as well as to PPPoE (IPCP, DHCPv6, and SLAAC).

Figure 109 **Direct Address Assignment using LUDB/RADIUS**

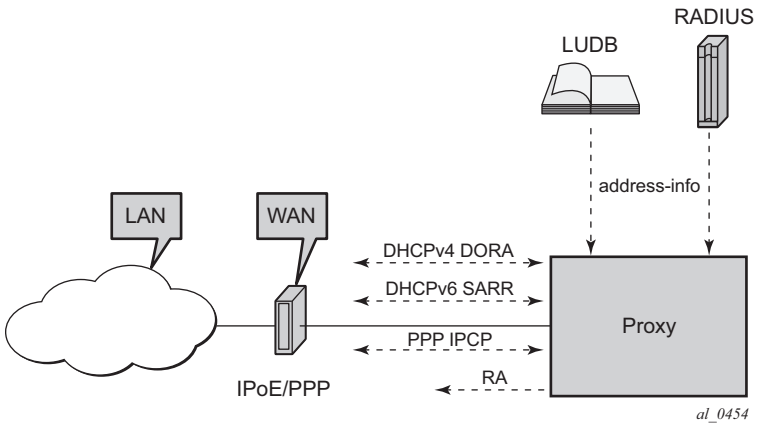

DHCP clients require a DHCP server address in the DHCP messages, and as such a *server-id* (IPv6) and/or an *emulated-server* (IPv4) must be configured, see the example below. IPCP and SLAAC do not require a DHCP server so these commands can be omitted.

```
configure
    service
        ies 1
            subscriber-interface "sub-int-1" create
                --- snip ---
                group-interface "grp-int-1-1" create
                    ipv6
                        dhcp6
                            --- snip ---
                            proxy-server
                                server-id duid-ll
                                client-applications dhcp ppp
                                no shutdown
                            exit
                        exit
                    exit
                    arp-populate
                    dhcp
                        proxy-server
                            emulated-server 10.1.1.254
                            no shutdown
                        exit
                        --- snip ---
                    exit
```

The emulated-server address must be a unique address in one of the subnets allowed on this subscriber interface.

Indirect Address Assignment using a DHCP server

In the example shown in Figure 110, the IPv4/IPv6 address/prefix is obtained from a DHCP server, using a pool-name as returned by LUDB/RADIUS.

The DHCP server (DHCPv4/DHCPv6) can be co-located with the BNG node (internal DHCP server) or can be external to the BNG node.

This scenario applies to both DHCPv4/DHCPv6 as well as to PPP. PPP relies on an internal DHCP client and communicates with the PPP client through IPCP.

Figure 110 **Indirect Address Assignment using a DHCP Server**

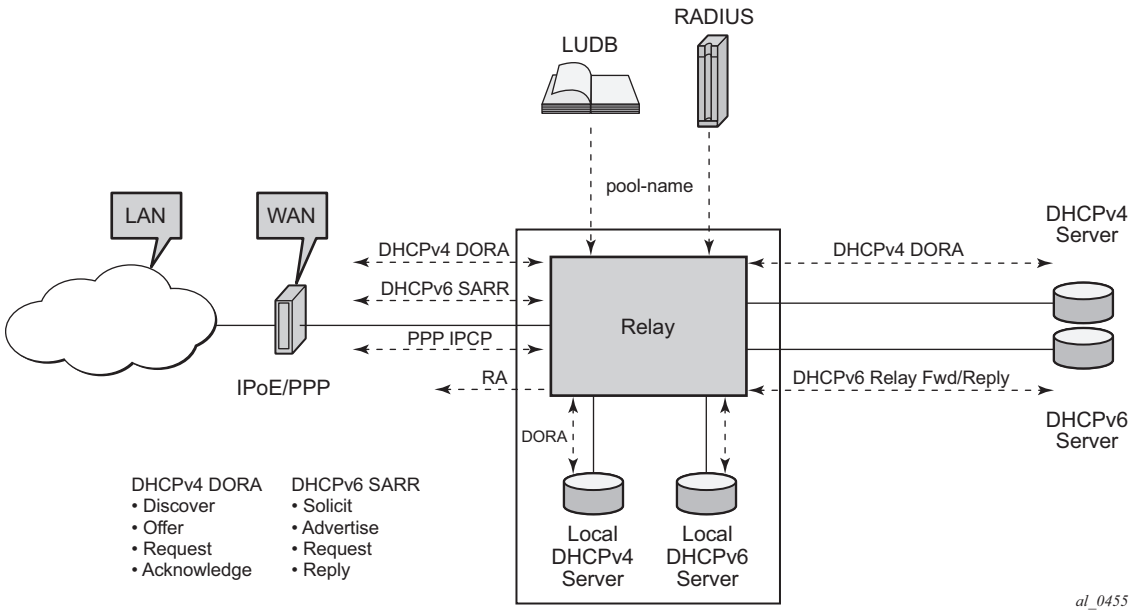

Relaying is configured for IPv4 and for IPv6 separately, at group-interface level.

The DHCPv4 relay agent is configured in the dhcp context:

- *gi-address* – The gateway IPv4-address used by the relay-agent.
- *server* - The DHCP server IPv4-address, 10.11.11.11 in the example.
- *client-applications dhcp ppp*
  - The DHCP server accepts requests for DHCP and for PPP.
- *option* – The options added/removed to/from messages towards the server. In the example the circuit-id, the remote-id and the pool-name are added.
- *trusted* – This parameter ensures that DHCP messages with option 82 included and the gi-address set to zero are being processed instead of being dropped.

The DHCPv6 relay agent is configured in the ipv6 dhcp6 relay context:

- *server* - The DHCP server IPv6 address, 2001:DB8::11 in the example.
- *client-applications dhcp ppp*

The DHCP server accepts requests for DHCP and for PPP.

```
configure
    service
        ies 1
            subscriber-interface "sub-int-1"
                --- snip ---
                group-interface "grp-int-1-1"
                    ipv6
                        --- snip ---
                        dhcp6
                            relay
                                server 2001:DB8::11
                                client-applications dhcp ppp
                                no shutdown
                            exit
                        exit
                    exit
                    dhcp
                        option
                            action replace
                            circuit-id
                            remote-id
                            vendor-specific-option
                                pool-name
                            exit
                        exit
                        server 10.11.11.1
                        trusted
                        lease-populate 100
                        client-applications dhcp ppp
                        gi-address 10.1.1.254
                        no shutdown
                    exit
```

Indirect Address Assignment using LAA

In the example shown in Figure 111 LUDB/RADIUS returns a pool-name which is resolved into an IP address using Local Address Assignment (LAA). LAA is implemented using a procedure call to an internal DHCP server, and not through the typical DHCPv4 DORA or DHCPv6 SARR sequence.

Figure 111 **Indirect Address Assignment using LAA**

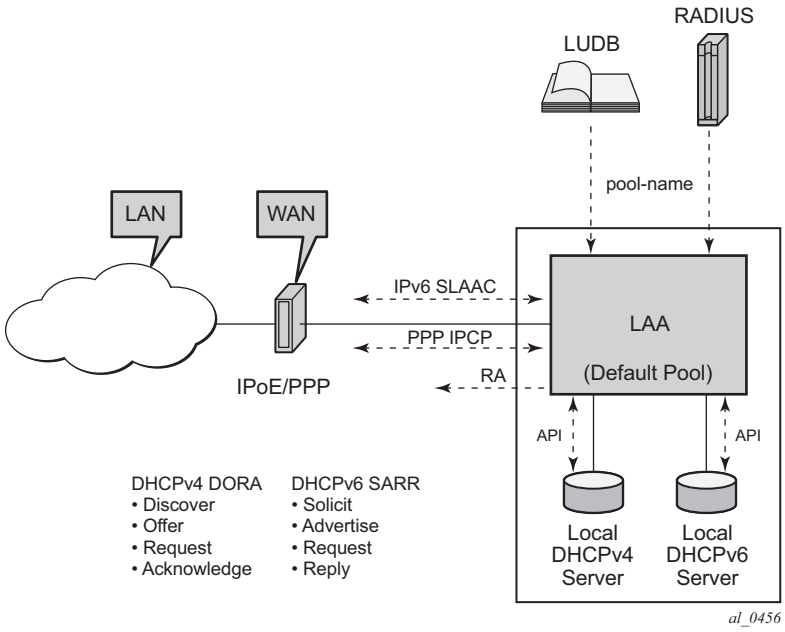

This scenario applies to PPPv4 hosts as well as to IPv6 SLAAC hosts. LAA can also be used to provide the IA-NA address in DHCPv6 proxy scenarios where an LUDB or a RADIUS server provides an IPv6 Delegated Prefix (Delegated-IPv6-Prefix) and an IPv6 WAN Address Pool (Framed-IPv6-Pool).

LAA is configured at group-interface level and overrules the relay scenario for the IPCP DHCP client (PPP) when both are configured.

The IPv4 parameters are configured in the local-address-assignment context:

- *server* – The internal DHCPv4 server to contact, identified by name.
- *default-pool* – Defines the pool name to use in case neither the LUDB nor the RADIUS server provides a pool-name.
- *client-application* – ppp-v4.

The IPv6 parameters are configured in the local-address-assignment ipv6 context:

- *server* – the DHCPv6 server to contact, identified by name.
- *client-application* – ppp-slaac and/or ipoe-wan.

An example is show below.

```
configure
    service
        ies 1
```

```
subscriber-interface "sub-int-1" create
    --- snip ---
    group-interface "grp-int-1-1" create
        local-address-assignment
            server "int-dhcp-v4"
            client-application ppp-v4
            default-pool "pool4-2"
            ipv6
                server "int-dhcp-v6"
                client-application ppp-slaac ipoe-wan
            exit
            no shutdown
        exit
```

LAA details are out of the scope of this example.

ESM Data Retrieval Examples

The following ESM data source scenarios are examined below:

- static host ESM data retrieval
- dynamic host ESM data retrieval:
  - RADIUS
  - LUDB + RADIUS access

DHCPv4 ACK python details are out of the scope of this example.

ESM Data Retrieval for Static Hosts

The example below shows the creation of a static host. The IP address is mandatory, a MAC address is optional.

```
configure
   service
       ies 1
           subscriber-interface "sub-int-1"
               group-interface "grp-int-1-1"
                   sap 1/1/1:111
                       description "sap for customer 1"
                       static-host ip 10.1.1.101 create
                           sla-profile "sla-prof-2"
                           sub-profile "sub-prof-2"
                           subscriber 03-7654321
                           no shutdown
                       exit
                   exit
```

Identification is not needed for static hosts: the subscriber-id, sub-profile and sla-profile used, are configured explicitly. Instantiation takes place at host creation time (**no shutdown**).

The following command shows the details for a single static host. The forwarding state is *Not Fwding*, meaning that no traffic can be forwarded to and from that host in this state.

```
 *A:BNG# show service id 1 static-host ip-address 10.1.1.101 detail
===============================================================================
Static Hosts for service 1
===============================================================================
Sap                          IP Address      Configured MAC    Dynamic MAC
  Subscriber                                   Admin State       Fwding State
-------------------------------------------------------------------------------
1/1/1:111                    10.1.1.101      N/A               N/A
  03-7654321                                   Up                Not Fwding
-------------------------------------------------------------------------------
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1-1
Sub Profile          : sub-prof-2
SLA Profile          : sla-prof-2
App Profile          : N/A
-------------------------------------------------------------------------------
Number of static hosts : 1
===============================================================================
*A:BNG# #
```

Once the MAC address of the host is learned through the ARP protocol, the forwarding state is set to *Fwding* and traffic now can be forwarded to and from that host.

```
*A:BNG# show service id 1 static-host
===============================================================================
Static Hosts for service 1
===============================================================================
Sap                          IP Address      Configured MAC    Dynamic MAC
  Subscriber                                   Admin State       Fwding State
-------------------------------------------------------------------------------
1/1/1:111                    10.1.1.101      N/A               00:00:00:99:99:99
  03-7654321                                   Up                Fwding
-------------------------------------------------------------------------------
Number of static hosts : 1
===============================================================================
*A:BNG# #
```

ESM Data Retrieval through RADIUS

A RADIUS server can provide ESM data, including:

- subscriber-id
- address information
- sub-profile-string
- sla-profile-string
- etc.

The BNG sends a RADIUS Access-Request including the User-Name attribute identifying the host for authentication purposes. The User-Name format is configurable on the BNG:

- MAC-address
- circuit-ID, remote-ID
- user@domain
- etc.

An excerpt from the Free-RADIUS user file used for this example is shown below. The first line of a block contains the User-Name with the credentials.

```
 00:00:00:22:22:22    Auth-Type := Local, Cleartext-Password := ""
        Alc-Subsc-ID-Str = "sub-22",
        Alc-Subsc-Prof-Str = "sub-string-1",
        Alc-SLA-Prof-Str = "sla-string-1"

BSAN64|40|1/1/2:300 Auth-Type := Local, User-Password == "LetMeIn"
        Alc-Subsc-ID-Str = "subscriber-300",
        Alc-Subsc-Prof-Str = "subpro1-string",
        Alc-SLA-Prof-Str = "slapro1-string",
        Session-Timeout = 600

sub202@provider  Cleartext-Password := "sub202"
        Alc-Subsc-ID-Str = "sub-44",
        Alc-Subsc-Prof-Str = "sub-string-3",
        Alc-SLA-Prof-Str = "sla-string-22",
        Framed-IP-Address = 10.2.1.202,
        Framed-IP-Netmask = 255.255.255.0
```

The lines following the User-Name with the credentials lists the attributes the RADIUS server returns after successfully authenticating the user, and include:

- Alc-Subsc-ID-Str
- Alc-Subsc-Prof-Str
- Alc-SLA-Prof-Str
- Framed-IP-Address
- Framed-IP-Netmask
- etc.

Authentication Policy

An authentication policy is needed for retrieval of the data from a RADIUS server, indirectly indicating the RADIUS server(s) to contact.

The following steps are needed to create an authentication policy:

1. Define one or more RADIUS servers.
2. Define a RADIUS server policy.
3. Define an authentication policy.

Step 1. The example defines a single RADIUS server, indicating the name, the address and the secret to use.

```
configure
    router
        radius-server
            server "server-1" address 192.168.202.84 secret secret-1 create
        exit
```

Step 2. The radius-server-policy refers to the server defined above. In this example the Base router is used to reach the RADIUS server.

```
configure
    aaa
        radius-server-policy "rad-serv-pol-1" create
            description "Radius AAA server policy"
            servers
                router "Base"
                server 1 name "server-1"
            exit
        exit
```

Step 3. The authentication policy specifies to include the circuit-id and the remote-id attributes towards the RADIUS server as well as the radius-server-policy to use. Multiple authentication policies can be defined.

```
configure
    subscriber-mgmt
        authentication-policy "auth-pol-1" create
            pppoe-access-method pap-chap
            include-radius-attribute
                circuit-id
                remote-id
            exit
            radius-server-policy "rad-serv-pol-1"
```

The authentication policy is applied at a group-interface. The example below shows both group interfaces using the same authentication policy.

```
configure
 service
    ies 1
        subscriber-interface "sub-int-1"
            group-interface "grp-int-1-1"
                -- snip --
                authentication-policy "auth-pol-1"
                -- snip –
            exit
            group-interface "grp-int-1-2"
                -- snip --
                authentication-policy "auth-pol-1"
                -- snip –
            exit
```

DHCP Host

The following command shows the strings returned by the RADIUS server for the DHCPv4 host with MAC-address 00:00:00:22:22:22.

```
A:BNG# show service id 1 dhcp lease-state mac 00:00:00:22:22:22 detail
===============================================================================
DHCP lease states for service 1
===============================================================================
Service ID           : 1
IP Address           : 10.1.1.9
Client HW Address    : 00:00:00:22:22:22
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1-1
SAP                  : 1/1/1:112
Up Time              : 0d 00:03:30
Remaining Lease Time : 9d 23:56:30
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "sub-22"
Sub-Profile-String   : "sub-string-2"         # profile string before translation
SLA-Profile-String   : "sla-string-2"         # profile string before translation
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : ""
Category-Map-Name    : ""

--- snip ---

Lease-Time           : 10d 00:00:00
DHCP Server Addr     : 10.11.11.1
Radius User-Name     : "00:00:00:22:22:22"
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:BNG#
```

The strings returned by the RADIUS server are translated according the subscriber identification profile for retrieval of the actual profile data.

The actual profiles being used can be found using following command.

```
*A:BNG# show service id 1 subscriber-hosts mac 00:00:00:22:22:22 detail
============================================================
Subscriber Host table
============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID Origin       Fwding State
------------------------------------------------------------
1/1/1:112              sub-22
  10.1.1.6
    00:00:00:22:22:22    N/A       DHCP         Fwding
------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : grp-int-1-1
Sub Profile           : sub-prof-2    # profile name after translation
SLA Profile           : sla-prof-2    # profile name after translation
App Profile           : N/A
Egress Q-Group        : N/A

--- snip ---

------------------------------------------------------------
Number of subscriber hosts : 1
============================================================
*A:BNG#
```

PPP Host

The following command shows the strings returned by the RADIUS server for the PPP host with user-name sub202@provider.

```
*A:BNG# show service id 1 ppp session user-name sub202@provider detail
===============================================================================
PPP sessions for service 1
===============================================================================

User-Name            : sub202@provider

Description          : svc:1 sap:1/1/1:212 mac:00:00:00:44:44:44 sid:1
Up Time              : 0d 00:17:13
Type                 : oE
Termination          : local
IP/L2TP-Id/If-Id     : 10.2.1.202
MC-Standby           : No
Session Time Left    : N/A

--- snip ---

Subscriber           : "sub-44"
```

```
Sub-Profile-String   : "sub-string-3"  # profile string before translation
SLA-Profile-String   : "sla-string-22" # profile string before translation
ANCP-String          : ""
Int-Dest-Id          : ""
App-Profile-String   : ""
Category-Map-Name    : ""

--- snip ---

Radius Class         :
Radius User-Name     : sub202@provider
--- snip ---
-------------------------------------------------------------------------------
No. of sessions: 1
===============================================================================
*A:BNG# #
```

The actual profiles being used can be found using following command.

```
*A:BNG# show service id 1 subscriber-hosts mac 00:00:00:44:44:44 detail
===============================================================
Subscriber Host table
===============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID Origin       Fwding State
---------------------------------------------------------------
1/1/1:212              sub-44
  10.2.1.202
    00:00:00:44:44:44    1         IPCP         Fwding
---------------------------------------------------------------
Subscriber-interface  : sub-int-2
Group-interface       : grp-int-2-1
Sub Profile           : sub-prof-3   # profile name after translation
SLA Profile           : sla-prof-20  # profile name after translation
App Profile           : N/A

--- snip ---

---------------------------------------------------------------
Number of subscriber hosts : 1
===============================================================
*A:BNG#
```

For both the DHCPv4 host as well as the PPPv4 host, the output aligns with the data provided by the RADIUS server and the defined profile maps.

ESM Data Retrieval through a Local User Database

A Local User Database (LUDB) can provide ESM data for DHCP and PPP hosts, including:

- subscriber-id
- address information
- sub-profile string
- sla-profile string
- MSAP service-id
- retail service-id
- etc.

Retrieval of data in an LUDB requires matching criteria, which can be one of the following items, or a combination of the following items, including:

- MAC-address
- circuit-id/remote-id
- username
- SAP-id
- etc.

The example below defines an LUDB named *ludb-1* which matches DHCP hosts by means of MAC address, and PPP hosts by means of username.

```
configure
    subscriber-mgmt
        local-user-db "ludb-1" create
            dhcp
                match-list mac
                host "host-121" create
                    host-identification
                        mac 00:00:00:aa:aa:aa
                    exit
                    address 10.1.1.121
                    identification-strings 254 create
                        subscriber-id "sub-121"
                        sla-profile-string "sla-string-3"
                        sub-profile-string "sub-string-2"
                    exit
                    options
                        subnet-mask 255.255.255.0
                        default-router 10.1.1.254
                    exit
                    ipv6-wan-address-pool "pool6-2"
                    ipv6-delegated-prefix-pool "pool6-2"
                    no shutdown
                exit
                host "host-122" create
                    host-identification
                        mac 00:00:00:bb:bb:bb
                    exit
                    auth-policy "auth-pol-1"
                    no shutdown
                exit
```

```
            exit
            ppp
                match-list username
                host "host-123" create
                    host-identification
                        username "user@domain"
                    exit
                    address 10.1.1.123/32
                    password pap user
                    identification-strings 254 create
                        subscriber-id "sub-123"
                        sla-profile-string "sla-string-3"
                        sub-profile-string "sub-string-1"
                    exit
                    no shutdown
                exit
            exit
            no shutdown
        exit
```

Note:

- The entire LUDB can be disabled.
- Individual host entries can be disabled.

An LUDB can be applied at group interface level in different contexts:

- dhcp
- ipv6 dhcp6
- ppp
- pppoe

For the LUDB to provide ESM data, no authentication policy may be applied at group interface level, and the LUDB itself must be in the no shutdown state.

The following example shows the LUDB named *ludb-1* applied to group-interface *grp-int-1-2* in the dhcp and the pppoe context.

```
configure
   service
       ies 1
           subscriber-interface "sub-int-1"
               address 10.1.1.254/24
               address 10.1.2.254/24
               group-interface "grp-int-1-2"
                   dhcp
                       --- snip ---
                       gi-address 10.1.1.254
                       user-db "ludb-1"
                       no shutdown
```

```
            exit
            no authentication-policy
            sap 1/1/1:121 create
                sub-sla-mgmt
                    def-sub-profile "sub-prof-1"
                    def-sla-profile "sla-prof-1"
                    sub-ident-policy "sub-id-pol-1"
                    multi-sub-sap
                    no shutdown
                exit
            exit
            sap 1/1/1:122 create
                sub-sla-mgmt
                    def-sub-profile "sub-prof-1"
                    def-sla-profile "sla-prof-1"
                    sub-ident-policy "sub-id-pol-1"
                    multi-sub-sap
                    no shutdown
                exit
            exit
            pppoe
                session-limit 100
                user-db "ludb-1"
                no shutdown
            exit
        exit
```

For *host-121*, with MAC-address 00:00:00:aa:aa:aa, all ESM data is provided by the LUDB. The detailed DHCP lease state shows the actual profile strings, the subscriber-ID and the IP address used. The Lease Info origin is set to UserDb.

```
A:BNG# show service id 1 dhcp lease-state mac 00:00:00:aa:aa:aa detail
===============================================================================
DHCP lease states for service 1
===============================================================================
Service ID           : 1
IP Address           : 10.1.1.121
Client HW Address    : 00:00:00:aa:aa:aa
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1-2
SAP                  : 1/1/1:121
--- snip ---
Sub-Ident            : "sub-121"
Sub-Profile-String   : "sub-string-2"
SLA-Profile-String   : "sla-string-3"
--- snip ---
Lease Info origin    : UserDb
--- snip ---
Ip-Netmask           : 255.255.255.0
Broadcast-Ip-Addr    : 10.1.1.255
Default-Router       : 10.1.1.254
Primary-Dns          : N/A
Secondary-Dns        : N/A
--- snip ---
Relay Agent Information
  Circuit Id         : BNG|1|grp-int-1-2|1/1/1:121
Radius User-Name     : ""
```

```
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
*A:BNG#
```

The detailed DHCPv6 lease state for the same host is shown below. The Lease State info now is DHCP for both IA-NA as well as for IA-PD where the IA-NA and the IA-PD are allocated by the DHCP server based on the ipv6-wan-address-pool and the ipv6-delegated-prefix-pool respectively, as indicated by the LUDB.

```
*A:BNG# show service id 1 dhcp6 lease-state mac 00:00:00:aa:aa:aa detail
===============================================================================
DHCP lease states for service 1
===============================================================================
Service ID           : 1
IP Address           : 2001:DB8:201::1/128
Client HW Address    : 00:00:00:aa:aa:aa
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1-2
SAP                  : 1/1/1:121
--- snip ---
Sub-Ident            : "sub-121"
Sub-Profile-String   : "sub-string-2"
SLA-Profile-String   : "sla-string-3"
--- snip ---
Pool Name            : "pool6-2"
Dhcp6 Server Addr    : 2001:DB8::11
--- snip ---
Lease Info origin    : DHCP
--- snip ---
Radius User-Name     : ""
-------------------------------------------------------------------------------
Service ID           : 1
IP Address           : 2001:DB8:202::/56
Client HW Address    : 00:00:00:aa:aa:aa
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1-2
SAP                  : 1/1/1:121
--- snip ---
Sub-Ident            : "sub-121"
Sub-Profile-String   : "sub-string-2"
SLA-Profile-String   : "sla-string-3"
--- snip ---
Pool Name            : "pool6-2"
Dhcp6 Server Addr    : 2001:DB8::11
--- snip ---
Lease Info origin    : DHCP
--- snip ---
Radius User-Name     : ""
-------------------------------------------------------------------------------
Number of lease states : 2
===============================================================================
*A:BNG#
```

Because only an authentication policy is defined for *host-122*, with MAC address 00:00:00:bb:bb:bb, the profile strings and the subscriber-ID are provided by the RADIUS server. The IP address is provided by the DHCP server. The Lease Info origin is set to DHCP.

```
A:BNG# show service id 1 dhcp lease-state mac 00:00:00:bb:bb:bb detail
===============================================================================
DHCP lease states for service 1
===============================================================================
Service ID           : 1
IP Address           : 10.1.1.5
Client HW Address    : 00:00:00:bb:bb:bb
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1-2
SAP                  : 1/1/1:122
--- snip ---
Sub-Ident            : "sub-122"
Sub-Profile-String   : "sub-string-3"
SLA-Profile-String   : "sla-string-1"
--- snip ---
Lease Info origin    : DHCP

Ip-Netmask           : 255.255.255.0
Broadcast-Ip-Addr    : N/A
Default-Router       : 10.1.1.254
--- snip ---
Lease-Time           : 10d 00:00:00
DHCP Server Addr     : 10.11.11.1

Relay Agent Information
  Circuit Id         : BNG|1|grp-int-1-2|1/1/1:122
Radius User-Name     : "00:00:00:bb:bb:bb"
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:BNG#
```

For *host-123*, the subscriber-ID, the profile strings and the IP-address are provided by the LUDB. Similar to *host-121*, no interaction with a RADIUS server is needed at all. The IP origin is set to local-user-db.

```
A:BNG# show service id 1 ppp session user-name "user@domain" detail
===============================================================================
PPP sessions for service 1
===============================================================================

User-Name            : user@domain

Description          : svc:1 sap:1/1/1:122 mac:00:00:00:cc:cc:cc sid:1
Up Time              : 0d 00:46:10
Type                 : oE
-- snip --
PPP MTU              : 1492
PPP Auth-Protocol    : PAP
PPP User-Name        : user@domain
```

```
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1-2

IP Origin            : local-user-db
DNS Origin           : none
NBNS Origin          : none

Subscriber           : "sub-123"
Sub-Profile-String   : "sub-string-1"
SLA-Profile-String   : "sla-string-3"
-- snip --
IP Address           : 10.1.1.123/32
-- snip --
Circuit-Id           :
Remote-Id            :

Radius Session-TO    : N/A
Radius Class         :
Radius User-Name     :
Logical-Line-Id      :
-------------------------------------------------------------------------------
No. of sessions: 1
===============================================================================
A:BNG#
```

LUDB details are out of the scope of this example.

Optional ESM Data

Limits

The maximum number of hosts, subscribers, sessions and leases are checked during the host instantiation process, and can be defined at following levels:

- sla-profile level: host-limit
- sap-level: multi-sub-sap
- group-interface level: lease-populate, session-limit, sap-session-limit

The sla-profile optionally defines a maximum number of hosts allowed by this profile. An example is shown below.

```
configure
  subscriber-mgmt
       sla-profile "sla-prof-3" create
           host-limit 4
           ingress
               qos 100
```

```
                exit
            exit
            egress
                qos 100
                exit
            exit
        exit
```

By default only hosts from a single subscriber can connect through a SAP. This condition can be lifted as follows.

```
configure
    service
        ies 1
            subscriber-interface "sub-int-1"
                group-interface "grp-int-1-1"sub-sla-mgmt
                    sap 1/1/1:112 create
                        sub-sla-mgmt
                            def-sub-profile "sub-prof-1"
                            def-sla-profile "sla-prof-1"
                            sub-ident-policy "sub-id-pol-2"
                            multi-sub-sap
                            no shutdown
                        exit
                    exit
```

The maximum number of leases assigned to subscriber hosts by the DHCP server can be defined at group interface level as follows.

```
configure
    service
        ies 1
            subscriber-interface "sub-int-1"
                address 10.1.1.254/24
                group-interface "grp-int-1-1"
                    dhcp
                        proxy-server
                            emulated-server 10.1.1.254
                            no shutdown
                        exit
                        server 10.11.11.1
                        trusted
                        lease-populate 100
                        client-applications dhcp ppp
                        gi-address 10.1.1.254
                        no shutdown
                    exit
```

The maximum number of sessions for PPP as well as for PPPoE, including the sap-session-limit, is also defined at group interface level, as the example below shows.

```
 configure
    service
        ies 1
            subscriber-interface "sub-int-1"
                group-interface "grp-int-1-1"
```

```
ppp
    session-limit 50
exit
pppoe
    session-limit 40
    sap-session-limit 30
    no shutdown
exit
```

Filtering

Filter policies allow for selectively dropping, forwarding or redirecting ingress/egress traffic. They are also known as access control lists (ACLs) and can optionally be included in the SLA-profile.

An example sla-profile with IP-filters is show below.

```
configure
    subscriber-mgmt
        sla-profile "sla-prof-3" create
            host-limit 4
            ingress
                qos 100
                exit
                ip-filter 1
                ipv6-filter 1
            exit
            egress
                qos 100
                exit
                ip-filter 2
                ipv6-filter 2
            exit
```

Filter policies must be defined before they can be used in an SLA profile.

Accounting

Accounting policies define how to count the traffic for billing purposes on a per service basis. Two types of accounting are available:

- RADIUS accounting
- XML accounting

For RADIUS accounting, the accounting data is stored on the RADIUS accounting server. For XML accounting, the accounting data is stored locally on a flash-disk on the node.

The example below shows the definition of a radius-accounting-policy and reuses the radius-server-policy defined before.

```
configure
    subscriber-mgmt
        radius-accounting-policy "rad-acct-pol-1" create
            host-accounting interim-update
            update-interval 5
            include-radius-attribute
                framed-ip-addr
                sla-profile
                sub-profile
            exit
            radius-accounting-server
                router "Base"
                server 1 address 192.168.202.84 secret secret-1
            exit
            radius-server-policy "rad-serv-pol-1"
        exit
```

The radius-accounting policy is referred to from the SUB profile as shown in the example below.

```
configure
    subscriber-mgmt
        sub-profile "sub-prof-1" create
            radius-accounting-policy "rad-acct-pol-1"
            ingress
                policer-control-policy "pol-ctrl-1"
                exit
            exit
            egress
                scheduler-policy "down-1"
                exit
            exit
        exit
```

For XML accounting, an accounting policy is referred to from the SUB profile.

```
configure
    subscriber-mgmt
        sub-profile "sub-prof-2" create
            accounting-policy 10
            collect-stats
            ingress
                scheduler-policy "sched-up-1"
                exit
            egress
                scheduler-policy "sched-down-1"
                exit
            exit
        exit
```

The **collect-stats** command activates the generation of accounting files.

Accounting policies must be defined before they can be used in a SUB profile.

RADIUS accounting and XML accounting details are out of the scope of this example.

Application Profile

An application profile is needed for supporting Application Assurance.

Traffic is diverted to an MS-ISA MDA which processes the traffic according the application profile, which for that purpose is assigned to:

- an ESM subscriber, or an group of ESM subscribers
- a SAP
- a spoke SDP

ESM subscribers are assigned an application profile either statically or dynamically.

The example below shows the assignment of an application profile to a static host.

```
configure
    service
        ies 1
            subscriber-interface "sub-int-1"
                group-interface "grp-int-1-1"
                    sap 1/1/1:111
                        description "sap for customer 1"
                        static-host ip 10.1.1.101 create
                            app-profile "app-prof-3"
                            sla-profile "sla-prof-2"
                            sub-profile "sub-prof-1"
                            subscriber 03-7654321
                            no shutdown
                        exit
                    exit
```

For dynamic hosts, the same rules apply as for the sub-profile scriber and the sla-profile, meaning that the app-profile can be found through:

1. LUDB
2. RADIUS
3. DHCPv4 ACK python
4. Default

The example below shows the assignment of a default application profile at the SAP level.

```
configure
    service
        ies 1
            subscriber-interface "sub-int-1"
                group-interface "grp-int-1-1"sub-sla-mgmt
                    sap 1/1/1:112 create
                        sub-sla-mgmt
                            def-sub-profile "sub-prof-1"
                            def-sla-profile "sla-prof-1"
                            def-app-profile "app-prof-1"
                            sub-ident-policy "sub-id-pol-2"
                            multi-sub-sap
                            no shutdown
                        exit
                    exit
```

Application profiles must be defined before they can be referenced.

Application Assurance and Application Assurance subscribers details are out of the scope of this example.

Intermediate Destinations

An intermediate destination is an aggregation point in the network and identified through the intermediate destination identity (inter-dest-id).

Most typically Access Nodes (ANs) are considered intermediate destinations.

The inter-dest-id is an optional per subscriber attribute.

The intermediate destination is used in the BNG for supporting QoS, as an example:

- By shaping the aggregate rate of all egress traffic of subscribers connected to an AN to prevent congestion of the link towards that AN.
- To ensure fairness amongst the subscriber hosts connected to that AN. For that purpose the inter-dest-id is linked to an hierarchical scheduler via a virtual port (Vport).

The example below shows the assignment of the inter-dest-id to a static host.

```
configure
    service
        ies 1
            subscriber-interface "sub-int-1"
                group-interface "grp-int-1-1"
                    sap 1/1/1:111
                        description "sap for customer 1"
                        static-host ip 10.1.1.101 create
                            inter-dest-id "bsan-1"
                            sla-profile "sla-prof-2"
```

```
            sub-profile "sub-prof-1"
            subscriber 03-7654321
            no shutdown
        exit
    exit
```

The inter-dest-id does not to be translated.

Intermediate destination details are is out of the scope of this example.

Conclusion

This chapter explains basic ESM concepts including the definition of subscribers and subscriber hosts. The host identification and instantiation process was explained, indicating the mandatory and optional ESM data fetched during the process (address information, sub-profile, sla-profile, app-profile, inter-dest-id, etc.). The address information retrieval process was explained. The different sources from where the ESM data can originate were mentioned, including some examples. Also the sub-ident-policy and the authentication-policy, needed during the identification and instantiation process, are explained in detail.

ESM IPv4: Multicast in a Wholesale/Retail Scenario

This chapter describes ESM IPv4 multicast configurations in a wholesale/retail scenario.

Topics in this chapter include:

Applicability

This configuration example is applicable to the 7750 SR-7/12/12e with IOM3-XP and IMMs, the 7450 ESS -7/12 chassis in mixed mode with IOM3-XP and IMMs, and also to the 7750 SR-c4/12 platforms, and requires chassis mode C as a minimum. Note that the 7450 will only operate as an L2TP Access Concentrator (LAC) for L2TP services.

The configuration was tested on release 11.0.R1 and covers both IPoE and PPPoE subscribers.

Overview

Triple Play Service Delivery Architecture (TPSDA) allows operators to integrate High Speed Internet (HSI), voice and video services within a single network infrastructure. The goal of this configuration example is to provide a walk through of a wholesale/ retail multicast setup.

There are two wholesale/retail models in TPSDA. In the first model, the retail service is co-located with the wholesale service whereas in the second model, for PPP services only, the retail service is on a separate BNG. The network topology shown in Figure 112 is the first model. It consists of two 7750s; BNG-1 is a wholesaler Broadband Network Gateway (BNG) with the retail service co-located and the second is a retailer router. Figure 113 shows the second model where the retail service is a separate router and the connection between the wholesale and retail utilizes L2TP. The 7450 in both cases is used as an aggregation switch to aggregate all subscribers.

Figure 112 **Wholesale/Retail Model 1**

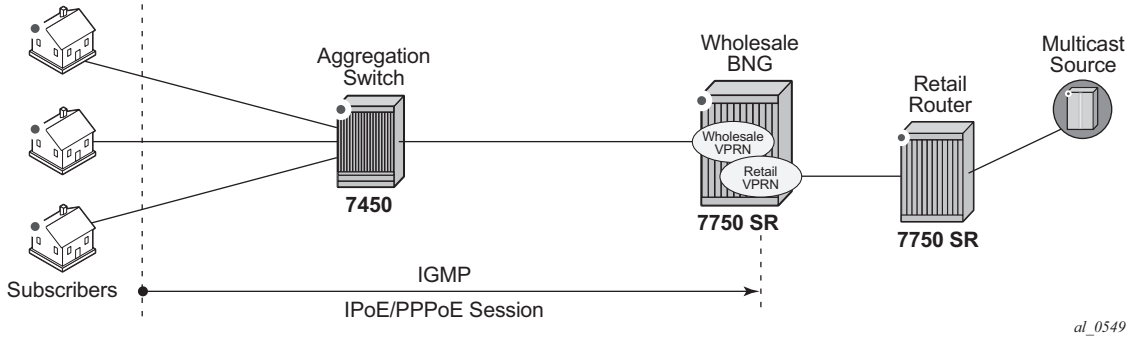

Figure 113 **Wholesale/Retail Model 2**

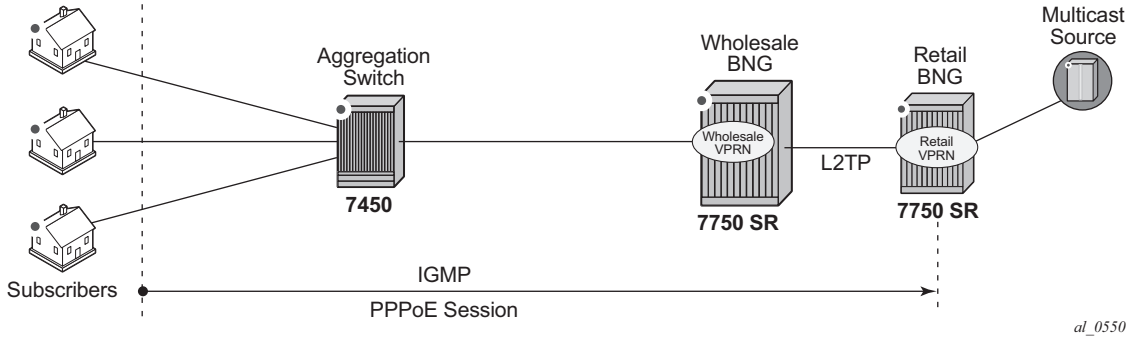

The second 7750 is connected directly to the multicast source. On the access side, the BNG is connected to an aggregation switch aggregating both PPPoE and IPoE subscribers.

There are two basic requirements for a subscriber to receive multicast streams. First, the group interface for the subscribers must have IGMP enabled. Second, the Enhanced Subscriber Management (ESM) subscriber must be allowed to receive multicast streams by having IGMP enabled. When both requirements are met, the BNG will process the subscribers' IGMP messages, otherwise, IGMP messages are dropped. All customer premise device (CPE) originated IGMP messages are aggregated via the 7450 and passed onto the wholesale BNG. It is always the retail VPRN that processes the IGMP messages. The wholesale VPRN SAPs performs the forwarding of the actual multicast streams.

Configuration

Note that a basic knowledge of multicast and ESM is assumed.

ESM Wholesale-Retail Multicast

There are various ways to provide wholesale and retail multicast function.

- For the IPoE and PPPoE Layer 3 wholesale/retail model, the wholesale and the retail services reside on separate VPRNs.
- For the PPPoE Layer 2 wholesale/retail model, L2TP is used.

ESM Layer 3 Wholesale-Retail Multicast

Figure 114 depicts a Layer 3 wholesale/retail scenario for both IPoE and PPPoE. The first BNG contains both the wholesale and retail configuration. There are two options for the retail BNG to deliver the multicast streams to the wholesale BNG:

1. MVPN between the BNGs

   or
2. If using a routed interface between the BNGs, multicast routing is required.

This example will use the second option for delivery of the multicast streams in order to keep the configuration simple.

Figure 114 **Layer 3 Wholesale/Retail**

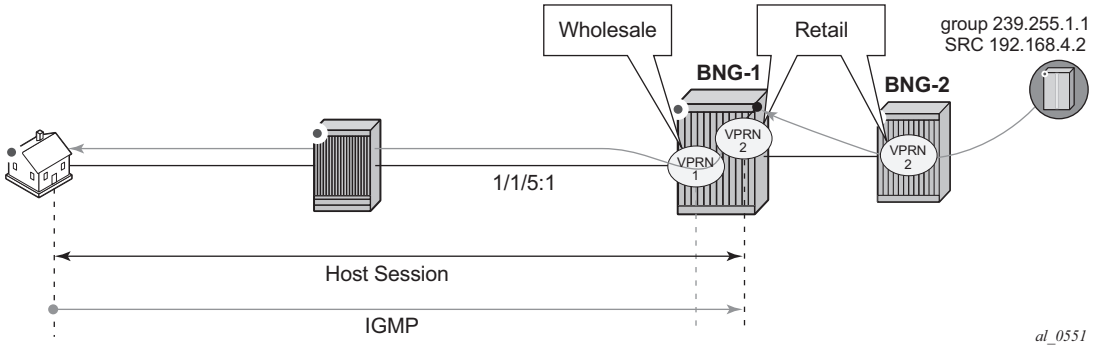

Step 1. Below is a configuration extract from the wholesale service on BNG-1 with the group interface added to IGMP. This configuration applies to both IPoE and PPPoE.

```
*A:BNG-1> config>service>vprn>sub-if# info
----------------------------------------------
                unnumbered "system"
                group-interface "wholesale-sub-int-1" create
                    dhcp
                        client-applications dhcp ppp
                        no shutdown
                    exit
                    authentication-policy "auth-policy-1"
                    sap 1/1/5:1 create
                        sub-sla-mgmt
                            def-sub-id use-sap-id
                            def-sub-profile "multicast-profile-1"
                            def-sla-profile "sla-profile-1"
                            sub-ident-policy "sub-ident-policy-1"
                            multi-sub-sap 10
                            no shutdown
                        exit
                    exit
                    pppoe
                        session-limit 10
                        sap-session-limit 10
                        no shutdown
                    exit
                exit
                igmp
                    group-interface "wholesale-group-int-1"
                        no shutdown
                    exit
                    no shutdown
                exit
```

Step 2. Also on BNG-1, a separate VPRN is configured for the retailer. The retail configuration is a little different from the wholesale configuration. Below is a configuration extract from the retail VPRN with IGMP enabled. This configuration is applicable to both IPoE and PPPoE. The multicast streams received in the retail VPRN are forwarded to the wholesale VPRN. Other retail VPRNs can offer multicast streams as well, and the same multicast addresses can be re-used as long as the address is assigned to a different retail VPRN.

```
*A:BNG-1> config>service>vprn# info
----------------------------------------------
            route-distinguisher 65536:2
            subscriber-interface "retail-sub-int-1" fwd-service 1
                                 fwd-subscriber-interface "wholesale-sub-int-
1" create
                address 10.255.255.254/8
                dhcp
                    server 192.168.0.1
                    client-applications dhcp ppp
                    gi-address 10.255.255.254
                    lease-populate 10
                    no shutdown
                exit
            exit
            igmp
                group-interface fwd-service 1 "wholesale-group-int-1
                    no shutdown
                exit
            exit
            ospf 192.168.2.2
                area 0.0.0.0
                    interface "system"
                        no shutdown
                    exit
                    interface "retail-sub-int-1"
                        no shutdown
                    exit
                    interface "int-BNG-1-BNG-2"
                        no shutdown
                    exit
                exit
            exit
            pim
                interface "int-BNG-1-BNG-2"
                exit
            exit
```

Step 3. Per host replication is mandatory in a wholesale/retail scenario. A single wholesale SAP might be shared among different retailers. A wholesale host that has requested a multicast group will always have the multicast delivered directly. Other hosts on the SAPs might belong to a different retailer and therefore 1) retailers might not have the same multicast group and sources and 2) their bandwidth should not be impacted by other hosts' multicast. Per-host replication is configured in the **igmp-policy** *igmp-policy-1*. This is mandatory for both IPoE and PPPoE subscribers.

```
*A:BNG 1> config>subscr-mgmt>igmp-policy# info
----------------------------------------------
            per-host-replication
```

Step 4. The interfaces are added to OSPF and to PIM on the retail BNG that is connected to the multicast source.

```
*A:BNG-2> config>service>vprn# info
----------------------------------------------
            ospf
                area 0.0.0.0
                    interface "system"
                        no shutdown
                    exit
                    interface "int-BNG-2-BNG-1"
                        no shutdown
                    exit
                    interface "int-multicast-source"
                        no shutdown
                    exit
                exit
            exit
            pim
                interface "int-BNG-2-BNG-1"
                exit
                interface "int-multicast-source"
                exit
                rp
                    static
                        address 192.168.4.1
                            group-prefix 224.0.0.0/4
                        exit
                    exit
                exit
            exit
```

With the above the configuration, the wholesale/retail setup is ready to process IGMP messages. Now send an IGMPv3 request to the wholesale SAP. The (S,G) is (192.168.4.2, 239.255.1.1) and the subscriber IP address is 10.0.0.2. The output below shows that the (S,G) is not registered in the wholesale VPRN but is in the retail VPRN.

```
*A:BNG-1> show router 1 igmp group
===============================================================================
```

```
IGMP Interface Groups
===============================================================================
===============================================================================
IGMP Host Groups
===============================================================================
===============================================================================
IGMP SAP Groups
===============================================================================
No Matching Entries
===============================================================================


*A:BNG-1> show router 2 igmp group
===============================================================================
IGMP Interface Groups
===============================================================================
===============================================================================
IGMP Host Groups
===============================================================================
(192.168.4.2,239.255.1.1)
    Fwd List  : 10.0.0.2             Up Time : 0d 00:13:01
===============================================================================
IGMP SAP Groups
===============================================================================
-------------------------------------------------------------------------------
(*,G)/(S,G) Entries : 1
===============================================================================
```

To view all subscribers' (S,G) pairs, use the following command.

```
*A:BNG-1> show service active-subscribers igmp detail
===============================================================================
Active Subscribers Detail
===============================================================================
Subscriber                      IGMP-Policy
  HostAddr                      GrpItf                            NumGroups
    GrpAddr                     Type            Up-Time           Mode
      SrcAddr                   Type                              Blk/Fwd
-------------------------------------------------------------------------------
video_user_01                   igmp-policy-1
  10.0.0.2                      whole-sale                        1
    239.255.1.1                 Dynamic         0d 01:37:55       Include
      192.168.4.2               Dynamic                           Fwd
-------------------------------------------------------------------------------
Number of Subscribers : 1
===============================================================================
```

Only the retail VPRN is responsible for processing the IGMP messages. Therefore to troubleshoot a wholesale/retail setup, debug is only relevant on the retail router instance.

```
debug
    router "2"
        igmp
            group-interface fwd-service "1" "whole-sale"
```

```
            host "10.0.0.2"
            packet mode egr-ingr-and-dropped
        exit
    exit

7648 2013/05/24 16:59:41.02 EST MINOR: DEBUG #2001 vprn2 IGMP[14]
"IGMP[14]: RX-PKT
[013 07:56:53.680] IGMP host 10.0.0.2 V3 PDU: 10.0.0.2 -> 224.0.0.22 pduLen
20
    Type: V3 REPORT maxrespCode 0x0 checkSum 0xddf6
    Num Group Records: 1
        Group Record 0
        Type: ALW_NEW_SRCS, AuxDataLen 0, Num Sources 1
        Mcast Addr: 239.255.1.1
        Source Address List
            192.168.4.2

"

7649 2013/05/24 16:59:41.02 EST MINOR: DEBUG #2001 vprn2 IGMP[vprn2 inst 1
4]
"IGMP[vprn2 inst 14]: igmpIfGroupAdd
Adding 239.255.1.1 to IGMP host 10.0.0.2 database"

7650 2013/05/24 16:59:41.02 EST MINOR: DEBUG #2001 vprn2 IGMP[vprn2 inst 1
4]
"IGMP[vprn2 inst 14]: igmpProcessGroupRec
Process group rec ALW_NEW_SRCS received on host 10.0.0.2 for group 239.255.1.1 i
n mode INCLUDE. Num srcs 1"

7651 2013/05/24 16:59:41.02 EST MINOR: DEBUG #2001 vprn2 IGMP[vprn2 inst 1
4]
"IGMP[vprn2 inst 14]: igmpIfSrcAdd
Adding i/f source entry for host 10.0.0.2 (192.168.4.2,239.255.1.1) to IGMP fwdList
 Database, redir if N/A"
```

The same **debug** command can be used for troubleshooting IGMP leave messages as shown below.

```
7652 2013/05/24 16:59:43.90 EST MINOR: DEBUG #2001 vprn2 IGMP[14]
"IGMP[14]: RX-PKT
[013 07:56:56.560] IGMP host 10.0.0.2 V3 PDU: 10.0.0.2 -> 224.0.0.22 pduLen
20
    Type: V3 REPORT maxrespCode 0x0 checkSum 0xdcf6
    Num Group Records: 1
        Group Record 0
        Type: BLK_OLD_SRCS, AuxDataLen 0, Num Sources 1
        Mcast Addr: 239.255.1.1
        Source Address List
            192.168.4.2

"

7653 2013/05/24 16:59:43.90 EST MINOR: DEBUG #2001 vprn2 IGMP[vprn2 inst 1
4]
"IGMP[vprn2 inst 14]: igmpProcessGroupRec
Process group rec BLK_OLD_SRCS received on host 10.0.0.2 for group 239.255.1.1 i
n mode INCLUDE. Num srcs 1"
```

```
7654 2013/05/24 16:59:43.90 EST MINOR: DEBUG #2001 vprn2 IGMP[vprn2 inst 1
4]
"IGMP[vprn2 inst 14]: igmpProcessIfSrcTimerExp
Source Timer expired for IGMP host 10.0.0.2 (192.168.4.2,239.255.1.1)"

7655 2013/05/24 16:59:43.90 EST MINOR: DEBUG #2001 vprn2 IGMP[vprn2 inst 1
4]
"IGMP[vprn2 inst 14]: igmpIfSrcDel
Deleting i/f source entry for host 10.0.0.2 (192.168.4.2,239.255.1.1) from IGMP Dat
abase. DeleteFromAvl: 1 !Redir 0"

7656 2013/05/24 16:59:43.90 EST MINOR: DEBUG #2001 vprn2 IGMP[vprn2 inst 1
4]
"IGMP[vprn2 inst 14]: igmpIfGroupDel
Deleting 239.255.1.1 from IGMP host 10.0.0.2 database"
```

ESM L2TP Wholesale/Retail Multicast

As previously mentioned, the other option for PPPoE wholesale/retail is to use an L2TP connection as shown in Figure 115. LAC-1 contains the wholesale configuration while LNS-1 contains the retail configuration.

Figure 115 **L2TP Wholesale-Retail Multicast**

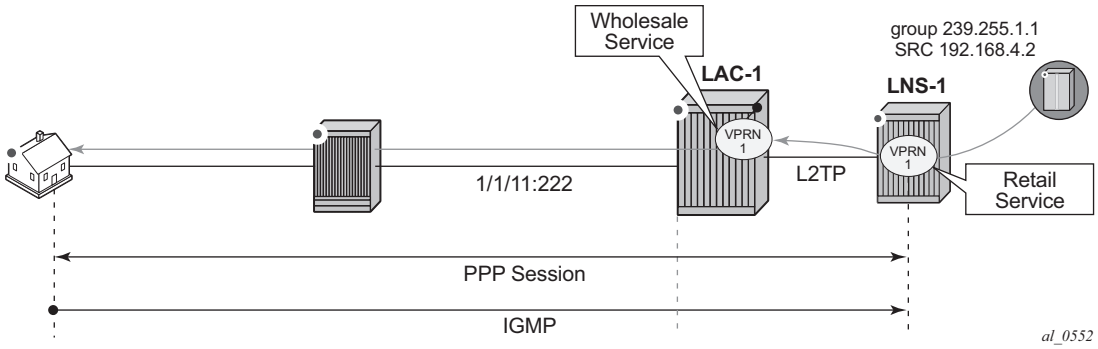

Below is a configuration extract for the wholesale L2TP Access Concentrator (LAC) in the wholesale service. It is using the local database, under the **pppoe user-db** configuration, to authenticate the subscriber. The wholesale LAC does not process any IGMP messages so it passes all messages to the retailer LNS.

```
*A:LAC-1> config>service>vprn>sub-if# info
----------------------------------------------
                description "L2TP"
                unnumbered "system"
                group-interface "LAC-sub-int-1" create
                    sap 1/1/11:222 create
                        sub-sla-mgmt
                            def-sub-id use-sap-id
```

```
                        def-sub-profile "multicast-profile-1
                        def-sla-profile "sla-profile-1"
                        sub-ident-policy "sub-ident-policy-1"
                        multi-sub-sap 10
                        no shutdown
                    exit
                exit
                pppoe
                    session-limit 10
                    sap-session-limit 10
                    user-db "ppp-db-1"
                    no shutdown
                exit
            exit
            l2tp
                group "l2tp-group-1" create
                    tunnel "tunnel-1" create
                        auto-establish
                        local-name "LAC"
                        peer 192.0.2.3
                        no shutdown
                    exit
                    no shutdown
                exit
                no shutdown
```

The retailer BNG serves as the L2TP Network Server (LNS). Below is a configuration extract for the LNS. IGMP must be enabled on the ESM group-interface in the retail service.

```
*A:LNS-1> config>service>vprn>sub-if# info
----------------------------------------------
            address 10.255.255.254/8
            group-interface "LNS-group-int-1" lns create
                sap-parameters
                    sub-sla-mgmt
                        def-sub-id use-sap-id
                        def-sub-profile "multicast-profile-1
                        def-sla-profile "sla-profile-1"
                        sub-ident-policy "sub-ident-policy-1"
                        multi-sub-sap 10
                        no shutdown
                    exit
                exit
                dhcp
                    server 192.168.0.1
                    client-applications ppp
                    gi-address 10.255.255.254
                    lease-populate 10
                    no shutdown
                exit
            exit
            l2tp
                group "l2tp-group-1" create
                    tunnel "tunnel-1" create
                        lns-group 1
                        ppp
                            authentication-policy "auth-policy-1"
```

```
                                default-group-interface "LNS-group-int-1" service-
id 1
                                mtu 1500
                                proxy-authentication always
                                proxy-lcp always
                            exit
                            remote-name "LAC"
                            no shutdown
                        exit
                        no shutdown
                    exit
                    no shutdown
                igmp
                    group-interface "LNS-group-int-1"
                        no shutdown
                    exit
                    no shutdown
```

With the above configuration applied, the wholesale/retail multicast setup can be verified. Firstly, send an IGMP message from the subscriber, the example below uses IGMPv3. The (S,G) sent is (192.168.4.2, 239.255.1.1) from the subscriber with IP address 10.0.0.2. The show commands below can be used to verify the multicast group being sent to the subscriber.

```
*A:LNS-1> show service active-subscribers igmp detail
===============================================================================
Active Subscribers Detail
===============================================================================
Subscriber                      IGMP-Policy
  HostAddr                      GrpItf                          NumGroups
    GrpAddr                     Type            Up-Time         Mode
      SrcAddr                   Type                            Blk/Fwd
-------------------------------------------------------------------------------
LNS1-pppoe-sub-01               igmp-policy-1
  10.0.0.2                      LNS                             1
    239.255.1.1                 Dynamic         0d 00:04:41     Include
      192.168.4.2               Dynamic                         Fwd
-------------------------------------------------------------------------------

Number of Subscribers : 1
===============================================================================
```

The IGMP group is not seen in the wholesale router instance (as shown by the first output below on LAC-1), however, it is seen in the retail router instance (as shown by the second output below on LNS-1).

```
*A:LAC-1> show router 1 igmp group
===============================================================================
IGMP Interface Groups
===============================================================================
===============================================================================
IGMP Host Groups
===============================================================================
===============================================================================
IGMP SAP Groups
```

```
===============================================================================
No Matching Entries
===============================================================================


*A:LNS-1> show router 1 igmp group
===============================================================================
IGMP Interface Groups
===============================================================================
===============================================================================
IGMP Host Groups
===============================================================================
(192.168.4.2,239.255.1.1)
    Fwd List  : 10.0.0.2         Up Time : 0d 00:08:27
===============================================================================
IGMP SAP Groups
===============================================================================
-------------------------------------------------------------------------------
(*,G)/(S,G) Entries : 1
===============================================================================
```

Only the retail BNG (LNS-1) is responsible for processing the IGMP messages. Therefore to troubleshoot ESM multicast for an L2TP service, the following debug commands are used on the LNS.

```
debug
    router "1"
        igmp
            group-interface "LNS-01"
            host "10.0.0.2"
            packet mode egr-ingr-and-dropped
        exit
    exit
7604 2013/05/24 16:55:49.46 EST MINOR: DEBUG #2001 vprn1 IGMP[8]

"IGMP[8]: RX-PKT
[013 07:53:02.120] IGMP host 10.0.0.2 V3 PDU: 10.0.0.2 -> 224.0.0.22
 pduLen 20
    Type: V3 REPORT maxrespCode 0x0 checkSum 0xddf6
    Num Group Records: 1
        Group Record 0
        Type: ALW_NEW_SRCS, AuxDataLen 0, Num Sources 1
        Mcast Addr: 239.255.1.1
        Source Address List
            192.168.4.2

"

7605 2013/05/24 16:55:49.46 EST MINOR: DEBUG #2001 vprn1 IGMP[vprn1 inst 8
]
"IGMP[vprn1 inst 8]: igmpIfGroupAdd
Adding 239.255.1.1 to IGMP host 10.0.0.2 database"

7606 2013/05/24 16:55:49.46 EST MINOR: DEBUG #2001 vprn1 IGMP[vprn1 inst 8
]
"IGMP[vprn1 inst 8]: igmpProcessGroupRec
Process group rec ALW_NEW_SRCS received on host 10.0.0.2 for group 239.255.1.1 in mo
de INCLUDE. Num srcs 1"
```

```
7607 2013/05/24 16:55:49.46 EST MINOR: DEBUG #2001 vprn1 IGMP[vprn1 inst 8
]
"IGMP[vprn1 inst 8]: igmpIfSrcAdd
Adding i/f source entry for host 10.0.0.2 (192.168.4.2,239.255.1.1) to IGMP fwd
List Database, redir if N/A"
```

The IGMP leave messages can also be seen in the debug, as shown below.

```
7615 2013/05/24 16:58:06.38 EST MINOR: DEBUG #2001 vprn1 IGMP[8]
"IGMP[8]: RX-PKT
[013 07:55:19.040] IGMP host 10.0.0.2 V3 PDU: 10.0.0.2 -> 224.0.0.22
 pduLen 20
    Type: V3 REPORT maxrespCode 0x0 checkSum 0xdcf6
    Num Group Records: 1
        Group Record 0
        Type: BLK_OLD_SRCS, AuxDataLen 0, Num Sources 1
        Mcast Addr: 239.255.1.1
        Source Address List
            192.168.4.2

"

7616 2013/05/24 16:58:06.38 EST MINOR: DEBUG #2001 vprn1 IGMP[vprn1 inst 8
]
"IGMP[vprn1 inst 8]: igmpProcessGroupRec
Process group rec BLK_OLD_SRCS received on host 10.0.0.2 for group 239.255.1.1 in mo
de INCLUDE. Num srcs 1"

7617 2013/05/24 16:58:06.38 EST MINOR: DEBUG #2001 vprn1 IGMP[vprn1 inst 8
]
"IGMP[vprn1 inst 8]: igmpProcessIfSrcTimerExp
Source Timer expired for IGMP host 10.0.0.2 (192.168.4.2,239.255.1.1)"

7618 2013/05/24 16:58:06.38 EST MINOR: DEBUG #2001 vprn1 IGMP[vprn1 inst 8
]
"IGMP[vprn1 inst 8]: igmpIfSrcDel
Deleting i/f source entry for host 10.0.0.2 (192.168.4.2,239.255.1.1) from IGMP
 Database. DeleteFromAvl: 1 !Redir 0"

7619 2013/05/24 16:58:06.38 EST MINOR: DEBUG #2001 vprn1 IGMP[vprn1 inst 8
]
"IGMP[vprn1 inst 8]: igmpIfGroupDel
Deleting 239.255.1.1 from IGMP host 10.0.0.2 database"
```

Conclusion

Multicast is an essential part of Triple Play Services. The SR/ESS TPSDA solution offering is much more than a baseline multicast delivery, it includes individual subscriber awareness and provides each retailer a separate routing context to manage their own multicast content. Subscriber awareness allows for the fine tuning of each subscriber multicast experience and also for troubleshooting on a per subscriber basis. This example provides a complete configuration walk through for multicast delivery for both IPoE and PPPoE in a wholesale/retail model.

ESM IPv4: Multicast with Redirection

This chapter describes ESM IPv4 multicast with redirection configurations.

Topics in this chapter include:

Applicability

This example is applicable to all 7750 SR-12 with IOM3-XP and IMMs, and needs chassis mode c as a minimum. This is also supported on 7450 ESS chassis in mixed mode and also on 7750 SR-c4/12 platform.

The configuration was tested on release 11.0R1 and covers both IPoE and PPPoE subscribers.

Overview

Triple Play Service Delivery Architecture (TPSDA) allows operators to integrate High Speed Internet (HSI), voice, and video services within a single network infrastructure. The goal of this example is to walk through a TPSDA multicast architecture with redirection. The topics are divided into the following sections:

- ESM (Enhanced Subscriber management) multicast baseline configuration
- Single BNG with redirection
- SRRP BNG configuration with static SAP
- IPoE ESM multicast configuration
- PPPoE ESM multicast configuration
- Subscriber Routed Redundancy Protocol (SRRP)
  - Multi-Chassis Synchronization (MCS) walk through

The network topology displayed in Figure 1 shows a typical TPSDA setup. It consists of three 7750s and a single 7450. Two 7750s are configured as Broadband Network Gateways (BNGs) and the third 7750 is configured as a P router. The 7450 is used as an aggregation switch to aggregate all subscribers. In ESM IPv4: Multicast with SRRP, multicast is directly distributed to a subscriber through a subscriber SAP. This example walks through another popular model which redirects all multicast streams to a common routed interface for all subscribers. When multicast is put on the common routed interface, one single copy of a multicast stream is delivered to multiple subscribers. In this model, per-subscriber replication of multicast streams is done on an access node or on the aggregation network in order to minimize the bandwidth consumed by the multicast traffic in access/aggregation.

Figure 116 **Network Topology Overview**

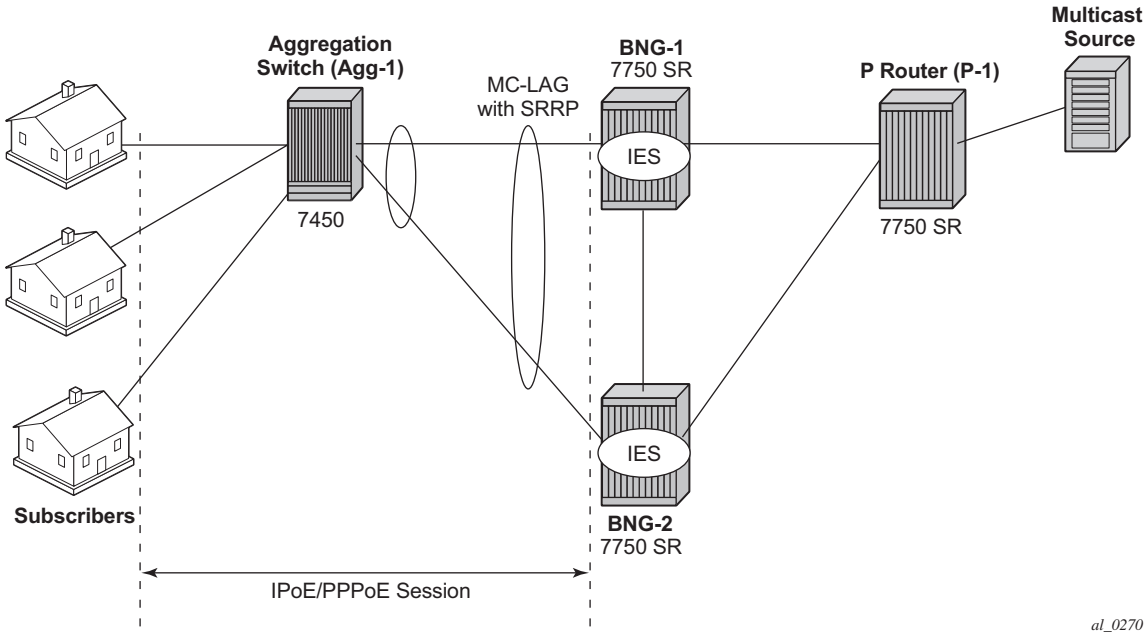

Figure 116 shows two BNGs configured with SRRP to provide redundancy. The P router is connected to the multicast source and is connected to both BNGs. The connections between the BNGs and the P router, and the multicast source and the P router, are also running PIM to provide multicast delivery. On the access side, the two BNGs are connected to an aggregation switch via MC-LAG aggregating the traffic for both PPPoE and IPoE subscribers. The BNGs facing the subscriber side are IGMP aware and will respond to any subscribers' IGMP requests.

There are two requirements for a subscriber to receive multicast streams. First, the ESM group-interface must have IGMP enabled. Second, the customization of each subscriber's subscriber profile to allow them to receive multicast streams. When both requirements are met, the BNG will process the subscribers' IGMP messages, otherwise, IGMP messages are simply dropped. All customer premise device (CPE)

IGMP messages are aggregated via the 7450 and passed onto the BNGs. Since the BNGs are running SRRP, the SRRP master is the only BNG processing and answering the IGMP messages. Protocol Independent Multicast (PIM) is then used between the BNG and the P router to request the multicast streams. If PIM is successful in retrieving the multicast group, the multicast stream is forwarded towards the individual subscribers. This is the typical multicast delivery for TPSDA.

Configuration

This example builds on the ESM multicast foundation discussed in ESM IPv4: Multicast with SRRP. It starts with a single BNG setup with redirection.

ESM Multicast Interface Redirection

Figure 117 shows a popular ESM multicast model that redirects all multicast streams to a dedicated router interface. When configuring a redirected interface be aware that:

1. Redirection between Global Routing Table (GRT) interfaces and VPRN interfaces is not supported
2. GRT interfaces are interfaces that reside in the base router or in an IES.
3. Redirection can be performed between interfaces in the GRT or between the interfaces in any VPRN (even different VPRNs).

The following steps start with a simple ESM multicast configuration for BNG, without redirection.

Figure 117 **Single BNG Setup with Multicast Redirection**

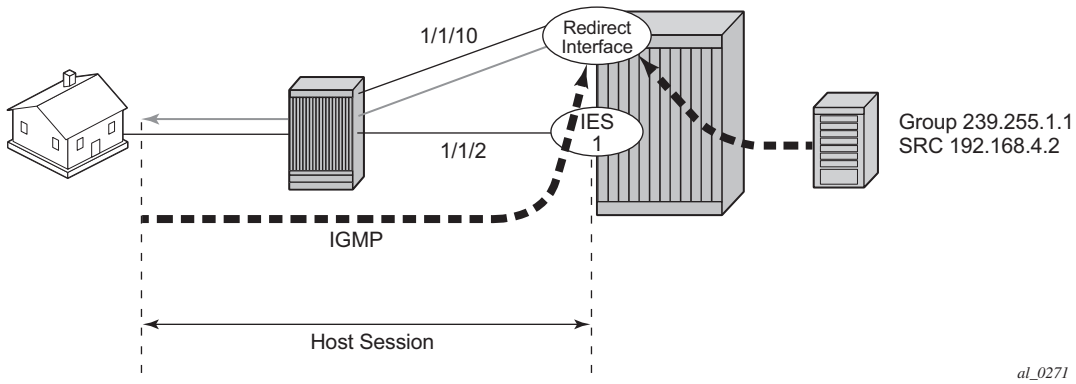

Step 1. Below is the BNG-1 configuration without multicast redirection. Subscribers are located in the 10.0.0.0/8 subnet. The multicast stream (S,G) is (192.168.4.2, 239.255.1.1). The local DHCP server is also on BNG-1.

```
*A:BNG-1>config>router>info
#------------------------------------------------
echo "Local DHCP Server Configuration"
#------------------------------------------------
        dhcp
            local-dhcp-server "dhcp-local-server" create
                use-gi-address scope pool
                pool "pool-1" create
                    subnet 10.0.0.0/8 create
                        options
                            subnet-mask 255.0.0.0
                            default-router 10.255.255.254
                        exit
                        address-range 10.0.0.10 10.0.0.254
                    exit
                exit
                no shutdown
            exit
        exit
#------------------------------------------------
echo "IP Configuration"
#------------------------------------------------
        interface "dhcp-lb1"
            address 192.168.0.1/32
            loopback
            local-dhcp-server “dhcp-local-server”
            no shutdown
        exit

*A:BNG-1>config>service>ies# info
----------------------------------------------
            description "BNG-1"
            interface “int-multicast-source" create
                address 192.168.4.1/30
                sap 1/1/2 create
                    no shutdown
                exit
            exit
            subscriber-interface "sub-int-1" create
                address 10.255.255.254/8
                group-interface "group-int-1" create
                    srrp-enabled-routing
                    dhcp
                        server 192.168.0.1
                        gi-address 10.255.255.254
                        lease-populate 10
                        no shutdown
                    exit
                    authentication-policy "auth-policy-1"
                    sap 1/1/5:4 create
                        sub-sla-mgmt
                            multi-sub-sap 10
                            no shutdown
```

```
                        exit
                    exit
                    pppoe
                        no shutdown
                    exit
                exit
            exit

*A:BNG-1>config>router# info
        interface "system"
            address 192.0.2.1/32
            no shutdown
        exit
        igmp
            group-interface “group-int-1”
                no shutdown
            exit
        exit
        pim
            interface "int-multicast-source"
            no shutdown
            rp
                static
                    address 192.0.2.1
                        group-prefix 224.0.0.0/4
                    exit
            exit
        exit

*A:BNG-1> config subscr-mgmt
     igmp-policy "igmp-policy-1" create
            exit
     exit all
     sub-profile "multicast-profile-1" create
            igmp-policy "igmp-policy-1"
     exit all
```

Step 2. Configure a router interface to redirect all multicast streams to, and then include the router interface in IGMP.

```
*A:BNG-1> config>service>ies# info
----------------------------------------------
            interface "redirected" create
                address 192.168.10.1/30
                sap 1/1/10 create
                exit
            exit

*A:BNG-1>config>router# info
            igmp
                interface "redirected"
----------------------------------------------
```

Step 3. Define a router redirection policy. This will redirect every (S,G) towards the redirected interface.

```
*A:BNG-1> config>router>policy-options# info
----------------------------------------------
            policy-statement "mcast_redirect_if"
                default-action accept
                    multicast-redirection fwd-service 1 "redirected"
                exit
            exit
```

Step 4. Apply the redirection policy created above in the igmp policy.

```
*A:BNG-1> config>subscr-mgmt>igmp-policy# info
----------------------------------------------
            redirection-policy "mcast_redirect_if"
----------------------------------------------
```

From this point on all multicast streams will be redirected to the “redirected” interface.

Now send an IGMPv3 join message and then use the *show router igmp group* command to verify that all multicast streams are redirected. In this example IGMPv3 is used with an (S,G) of (192.168.4.2, 239.255.1.1). The host has the IP address 10.0.0.10. Below is the output for PPPoE and for IPoE subscribers, shown separately.

a. Step 5a. Redirection with PPPoE subscribers: Viewing the multicast groups. In the PPPoE case, the multicast (S,G) shows up on both the redirected interface and the host.

```
*A:BNG-1> show router igmp group
===============================================================================
IGMP Interface Groups
===============================================================================

(192.168.4.2,239.255.1.1)                   Up Time : 0d 00:00:04
    Fwd List  : redirected
===============================================================================
IGMP Host Groups
===============================================================================
(192.168.4.2,239.255.1.1)
    Fwd List  : 10.0.0.10                   Up Time : 0d 00:00:04
===============================================================================
IGMP SAP Groups
===============================================================================
-------------------------------------------------------------------------------
(*,G)/(S,G) Entries : 2
===============================================================================
```

b. Step 5b. Redirection with IPoE subscribers: Viewing the multicast groups. In the IPoE case, the multicast (S,G) shows up on both the redirected interface and the SAP.

```
*A:BNG-1> show router igmp group
===============================================================================
IGMP Interface Groups
===============================================================================

(192.168.4.2,239.255.1.1)                  Up Time : 0d 00:00:04
    Fwd List  : redirected
===============================================================================
IGMP Host Groups
===============================================================================
===============================================================================
IGMP SAP Groups
===============================================================================
(192.168.4.2,239.255.1.1)
    Fwd List  : 10.0.0.10                  Up Time : 0d 00:00:04
-------------------------------------------------------------------------------
(*,G)/(S,G) Entries : 2
===============================================================================
```

Now the “redirected” interface is the only interface sending out multicast streams. The first command shows that the group interface does not register any multicast group (Num-Groups=0). The second command displays all multicast group are registered against the redirected interface (Num-Groups=1).

```
*A:BNG-1> show router igmp group-interface
===============================================================================
IGMP Group-Interfaces
===============================================================================
FwdSvc Group-Interface                 Adm/Opr-State        Import-Policy
       SAP                             Adm/Opr-Version      Num-Groups
-------------------------------------------------------------------------------
1      group-int-1                     Up/Up                none
       1/1/2                           3/3                  0
-------------------------------------------------------------------------------
Group-Interfaces = 1, SAPs = 1
===============================================================================

*A:BNG-1> show router igmp interface
===============================================================================
IGMP Interfaces
===============================================================================
Interface              Adm  Oper Querier         Cfg/Opr Num     Policy
                                                 Version Groups
-------------------------------------------------------------------------------
redirected             Up   Up   192.168.10.1      3/3     1       none
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
```

Debug facilities can be used to troubleshoot multicast redirection issues. The output below shows the multicast is redirected to a regular routed interface after an IGMP join.

```
7017 2013/05/24 09:27:50.65 EST MINOR: DEBUG #2001 ies1 IGMP[9]
"IGMP[9]: RX-PKT
[013 00:25:03.310] IGMP host 10.0.0.10 V3 PDU: 10.0.0.10 -> 224.0.0.22 pduLen
20
    Type: V3 REPORT maxrespCode 0x0 checkSum 0xddf6
    Num Group Records: 1
        Group Record 0
        Type: ALW_NEW_SRCS, AuxDataLen 0, Num Sources 1
        Mcast Addr: 239.255.1.1
        Source Address List
            192.168.4.2

"

7018 2013/05/24 09:27:50.65 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpIfGroupAdd
Adding 239.255.1.1 to IGMP host 10.0.0.10 database"

7019 2013/05/24 09:27:50.65 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpProcessGroupRec
Process group rec ALW_NEW_SRCS received on host 10.0.0.10 for group 239.255.1.1 i
n mode INCLUDE. Num srcs 1"

7020 2013/05/24 09:27:50.66 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpIfSrcAdd
Adding i/f source entry for host 10.0.0.10 (192.168.4.2,239.255.1.1) to IGMP fwdList
 Database, redir if interface redirected [ifIndex 13]"
```

The output below shows what happens when an IGMP leave message is sent so that the multicast stream is no longer being forwarded.

```
7024 2013/05/24 09:29:29.85 EST MINOR: DEBUG #2001 ies1 IGMP[9]
"IGMP[9]: RX-PKT
[013 00:26:42.510] IGMP host 10.0.0.10 V3 PDU: 10.0.0.10 -> 224.0.0.22 pduLen
20
    Type: V3 REPORT maxrespCode 0x0 checkSum 0xdcf6
    Num Group Records: 1
        Group Record 0
        Type: BLK_OLD_SRCS, AuxDataLen 0, Num Sources 1
        Mcast Addr: 239.255.1.1
        Source Address List
            192.168.4.2
"
7025 2013/05/24 09:29:29.85 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpProcessGroupRec
Process group rec BLK_OLD_SRCS received on host 10.0.0.10 for group 239.255.1.1 i
n mode INCLUDE. Num srcs 1"

7026 2013/05/24 09:29:29.85 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpProcessIfSrcTimerExp
```

```
Source Timer expired for IGMP host 10.0.0.10 (192.168.4.2,239.255.1.1)"

7027 2013/05/24 09:29:29.85 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpIfSrcDel
Deleting i/f source entry for host 10.0.0.10 (192.168.4.2,239.255.1.1) from IGMP Dat
abase. DeleteFromAvl: 1 Redir 0"

7028 2013/05/24 09:29:29.85 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpIfGroupDel
Deleting 239.255.1.1 from IGMP host 10.0.0.10 database"

7029 2013/05/24 09:29:29.85 EST MINOR: DEBUG #2001 ies1 IGMP MCS[9]
"IGMP MCS[9]: TX-MCS Data (GlblDel)
host 10.0.0.10
Key Type: HostGroup, Len: 13, Host : 10.0.0.10, Grp Addr: 239.255.1.1
Data Type: Group, Len: 16, Ver: 0, RecType: 1, Compat Mode: 3,
Num Fwd Srcs: 0, Num Blk Srcs: 0
"

7030 2013/05/24 09:29:29.85 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpMcsDelIfGroup
Deleting MCS entry for host 10.0.0.10, group 239.255.1.1, Glb"
```

ESM SRRP with MC-LAG

Figure 118 shows a numbered SRRP setup with MC-LAG SAPs serving both IPoE and PPPoE subscribers. ESM IPv4: Multicast with SRRP covers the configuration of regular SRRP SAPs, consequently this example provides configuration guidelines to use a different type of SAP: SRRP MC-LAG SAPs. Note that redirection on SRRP SAPs without MC-LAG is also supported. The configuration of the RADIUS server is out of the scope of this example.

Figure 118 **Network Topology with MC-LAG**

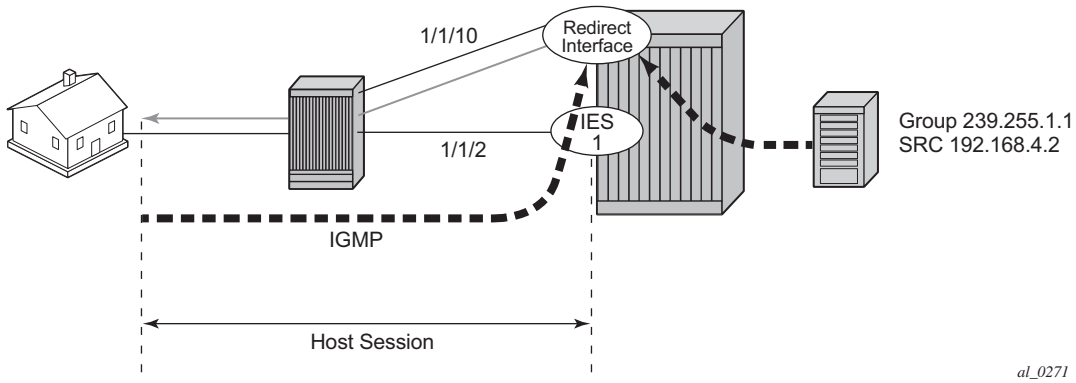

The baseline configuration for BNG-1 is shown below without any IGMP configuration. The configuration begins with the MC-LAG configuration. ESM is configured in an IES service but it is also possible to configure ESM in a VPRN. The redirection interface must be in the same routing instance as the group-interface, this applies to both regular SRRP SAPs and MC-LAG SAPs. In the following example, the MC-LAG is **lag-1**, customer data traffic is using VLAN 4, MC-LAG control traffic is using VLAN 5, and the redirected multicast streams are using VLAN 4094.

```
A:BNG-1>config>lag# info
----------------------------------------------
        mode access
        encap-type dot1q
        port 1/1/5 priority 1
        lacp active administrative-key 32768
        no shutdown

A:BNG-1>config>redundancy# info
----------------------------------------------
        multi-chassis
            peer 192.0.2.2 create
                mc-lag
                    lag 1 lacp-key 1 system-id 00:00:00:00:00:01 system-priority 100
                    no shutdown
                exit
                sync
                    igmp
                    srrp
                    sub-mgmt ipoe pppoe
                    port lag-1 create
                        range 4-4 sync-tag "mclagdata"
                        range 5-5 sync-tag "mclagcontrol"
                    exit
                    no shutdown
                exit
                no shutdown
            exit
        exit

A:BNG-1>config>service>ies# info
----------------------------------------------
            description "BNG-1"
            redundant-interface "MClink-BNG-1-BNG-2" create
                address 192.168.1.0/31
                ip-mtu 1500
                spoke-sdp 1:1 create
                    no shutdown
                exit
            exit
            interface "int-BNG-1-P-1" create
                address 192.168.2.1/30
                sap 1/1/2 create
                    no shutdown
                exit
            exit
            interface "lag-redirected" create
                address 192.168.10.253/24
                vrrp 1
```

```
                        backup 192.168.10.254
                    exit
                    sap lag-1:4094 create
                    exit
                exit
                subscriber-interface "sub-int-1" create
                    address 10.255.255.253/8 gw-ip-address 10.255.255.254 track-srrp 1
                    group-interface "group-int-1" create
                        dhcp
                            server 192.168.0.1
                            gi-address 10.255.255.253
                            lease-populate 10
                            no shutdown
                        exit
                        authentication-policy "auth-policy-1"
                        redundant-interface "MClink-BNG-1-BNG-2"
                        sap lag-1:1 create
                            sub-sla-mgmt
                                def-sub-id use-sap-id
                                def-sub-profile "multicast-profile-1"
                                def-sla-profile "sla-profile-1"
                                sub-ident-policy "sub-ident-policy-1"
                                multi-sub-sap 10
                                no shutdown
                            exit
                        exit
                        sap lag-1:5 create
                        exit
                        srrp 4 create
                            message-path lag-1:5
                            priority 200
                            no shutdown
                        exit
                        pppoe
                            no shutdown
                        exit
                    exit
                exit

*A:BNG-1>config>router# info
#------------------------------------------------
echo "IP Configuration"
#------------------------------------------------
        interface "int-BNG-1-BNG-2"
            address 192.168.6.1/30
            port 1/1/1:1
            no shutdown
        exit
        interface "system"
            address 192.0.2.1/32
            bfd 100 receive 100 multiplier 3
            no shutdown
        exit
        autonomous-system 65536
#------------------------------------------------
echo "OSPFv2 Configuration"
#------------------------------------------------
        ospf
            traffic-engineering
```

```
            area 0.0.0.0
                interface "system"
                    no shutdown
                exit
                interface "int-BNG-1-BNG-2"
                    interface-type point-to-point
                    metric 10000
                    no shutdown
                exit
                interface "sub-int-1"
                    no shutdown
                exit
                interface "int-BNG-1-P-1"
                    no shutdown
                exit
                interface "lag-redirected"
                    no shutdown
                exit
            exit
        exit
        pim
            interface “int-to_P_router”
        exit
```

The baseline configuration for BNG-2 is shown below without IGMP configuration. The default SRRP priority for BNG-2 is lower than the SRRP priority for BNG-1 and hence BNG-2 will be in standby mode.

```
A:BNG-2>config>lag# info
----------------------------------------------
        mode access
        encap-type dot1q
        port 1/1/5 priority 1
        lacp active administrative-key 32768
        no shutdown

A:BNG-2>config>redundancy# info
----------------------------------------------
        multi-chassis
            peer 192.0.2.1 create
                mc-lag
                    lag 1 lacp-key 1 system-id 00:00:00:00:00:01 system-priority 100
                    no shutdown
                exit
                sync
                    igmp
                    srrp
                    sub-mgmt ipoe pppoe
                    port lag-1 create
                        range 4-4 sync-tag "mclagdata"
                        range 5-5 sync-tag "mclagcontrol"
                    exit
                    no shutdown
                exit
                no shutdown
```

```
            exit
        exit


A:BNG-2>config>service>ies# info
----------------------------------------------
            description "BNG SRRP1"
            redundant-interface "MClink-BNG-1-BNG-2" create
                address 192.168.1.1/31
                ip-mtu 1500
                spoke-sdp 1:1 create
                    no shutdown
                exit
            exit
            interface "lag-redirected" create
                address 192.168.10.252/24
                vrrp 2
                    backup 192.168.10.254
                exit
                sap lag-1:4094 create
                exit
            exit
            interface "int-BNG-2-P-1" create
                address 192.168.3.1/30
                sap 1/1/2 create
                    no shutdown
                exit
            exit
            subscriber-interface "sub-int-1" create
                address 10.255.255.252/8 gw-ip-address 10.255.255.254 track-srrp 1
                group-interface "group-int-1" create
                    dhcp
                        server 192.168.0.1
                        lease-populate 10
                        gi-address 10.255.255.252
                        no shutdown
                    exit
                    authentication-policy "auth-policy-1"
                    redundant-interface "MClink-BNG-1-BNG-2"
                    sap lag-1:4 create
                        sub-sla-mgmt
                            def-sub-id use-sap-id
                            def-sub-profile "multicast-profile-1"
                            def-sla-profile "sla-profile-1"
                            sub-ident-policy "sub-ident-policy-1"
                            multi-sub-sap 10
                            no shutdown
                        exit
                    exit
                    sap lag-1:5 create
                    exit
                    srrp 1 create
                        message-path lag-1:5
                        no shutdown
                    exit
                    pppoe
                        no shutdown
                    exit
```

```
                exit
            exit

*A:BNG-2>config>router# info
#------------------------------------------------
echo "IP Configuration"
#------------------------------------------------
        interface "int-BNG-2-BNG-1"
            address 192.168.6.1/30
            port 1/1/1:1
            no shutdown
        exit
        interface "system"
            address 192.0.2.2/32
            bfd 100 receive 100 multiplier 3
            no shutdown
        exit
        autonomous-system 65536
#------------------------------------------------
echo "OSPFv2 Configuration"
#------------------------------------------------
        ospf
            traffic-engineering
            area 0.0.0.0
                interface "system"
                    no shutdown
                exit
                interface "int-BNG-2-BNG-1"
                    interface-type point-to-point
                    metric 10000
                    no shutdown
                exit
                interface "sub-int-1"
                    no shutdown
                exit
                interface "lag-redirected"
                    no shutdown
                exit
                interface "int-BNG-2-P-1"
                    no shutdown
                exit
            exit
        exit
        pim
            interface “int-BNG-2-P-1”
        exit
```

The baseline configuration for the 7450 aggregation switch is shown below. It has a LAG interface configured. There are two VPLSs. The first is VPLS 1 which is used to receive all redirected multicast traffic on VLAN 4094. The second is VPLS 2 which is responsible for passing all subscriber traffic on VLAN 4.

```
A:Agg-1>config>lag# info
----------------------------------------------
        mode access
```

```
        encap-type dot1q
        port 1/1/2
        port 1/1/3
        lacp active administrative-key 1
        no shutdown

*A:Agg-1>config>service>info
            vpls 1 customer 1 create
                sap lag-1:4094 create
                    no shutdown
                exit
                sap 1/1/1:4094 create
                    no shutdown
                exit
            no shutdown
            exit


*A:Agg-1>config>service>info
            vpls 2 customer 1 create
                sap lag-1:4 create
                    no shutdown
                exit
                sap 1/1/1:4 create
                    no shutdown
                exit
            no shutdown
            exit
```

The baseline configuration for the P router is shown below. It is now responsible for DHCP address assignment (moved from BNG-1 in the previous configuration to allow for redundant operations in case of failure of either BNG-1 or BNG-2) and is also attached to the multicast source.

```
*A:P-router>config>router>info
#------------------------------------------------
echo "Local DHCP Server Configuration"
#------------------------------------------------
        dhcp
            local-dhcp-server "dhcp-local-server" create
                use-gi-address scope pool
                pool "pool-01" create
                    subnet 10.0.0.0/8 create
                        options
                            subnet-mask 255.0.0.0
                            default-router 10.255.255.254
                        exit
                        address-range 10.0.0.10 10.0.0.254
                    exit
                exit
                no shutdown
            exit
        exit
#------------------------------------------------
echo "IP Configuration"
#------------------------------------------------
```

```
            interface "dhcp-lb1"
                address 192.168.0.1/32
                loopback
                local-dhcp-server "dhcp-local-server"
                no shutdown
            exit
            interface "int-P-1-BNG-1"
                address 192.168.2.2/30
                port 1/1/2
                no shutdown
            exit
            interface "int-P-1-BNG-2"
                address 192.168.3.2/30
                port 1/1/3
                no shutdown
            exit
            interface "P-1-multicast-source"
                address 192.168.4.1/30
                port 1/1/1
                no shutdown
            exit
            interface "system"
                address 192.0.2.3/32
                no shutdown
            exit

        #--------------------------------------------------
            ospf
                area 0.0.0.0
                    interface "system"
                        no shutdown
                    exit
                    interface "int-P-1-BNG-1"
                        no shutdown
                    exit
                    interface "int-P-1-BNG-2"
                        no shutdown
                    exit
                    interface "P-1-multicast-source"
                        no shutdown
                    exit
                exit
            exit
            pim
                interface "int-P-1-BNG-1"
                exit
                interface "int-P-1-BNG-2"
                exit
                interface "P-1-multicast-source"
                exit
            exit
```

Enable IGMP on Group Interface and Redirect Interface on the BNGs

The configuration below shows how to add the group-interface and redirect interface to IGMP. If ESM is configured in a VPRN, each VPRN will have its own IGMP instance. Remember to apply the following configuration to both BNG-1 and BNG-2.

```
*A:BNG-1>config>router>igmp# info
----------------------------------------------
                group-interface "group-int-1"
                    no shutdown
                exit
                interface "lag-redirected"
                    no shutdown
                exit
```

Next, the IGMP policy is configured to redirect all multicast to a dedicated interface. The following configuration outlines the steps necessary to enable multicast redirection.

Step 1. Define a router redirection policy. This will redirect every (S,G) towards the redirected interface.

```
*A:BNG-1> config>router>policy-options# info
----------------------------------------------
            policy-statement "mcast_redirect_if"
                default-action accept
                    multicast-redirection fwd-service 1 "lag-redirected"
exit
            exit
```

Step 2. Apply the redirection policy to the igmp-policy.

```
*A:BNG-1> config>subscr-mgmt>igmp-policy# info
----------------------------------------------
            redirection-policy "mcast_redirect_if"
----------------------------------------------
```

Step 3. Add multi-chassis synchronization of the redirected interface. This will synchronize the IGMP state on this MC-LAG interface.

```
*A:BNG-1>config>redundancy# info
----------------------------------------------
        multi-chassis
            peer 192.0.2.2 create
                sync
                    port lag-1 create
                        range 4-4 sync-tag "mclagdata"
                        range 5-5 sync-tag "mclagcontrol"
                        range 4094-4094 sync-tag "mclagmulticast"
```

```
igmp
srrp
sub-mgmt ipoe pppoe
exit
no shutdown
```

ESM IGMP IPoE walkthrough

With the baseline configuration applied, the BNG is ready to process IGMP messages and deliver multicast streams to the subscribers through the redirected interface. Figure 4 shows the message flow for IPoE subscribers requesting and receiving multicast traffic. The key points are highlighted in the dotted box:

- The group-interface and redirect interface must have IGMP enabled.
- The subscriber must be associated with an IGMP-policy via sub-profile.

Figure 119 **IPoE Multicast Message Flow**

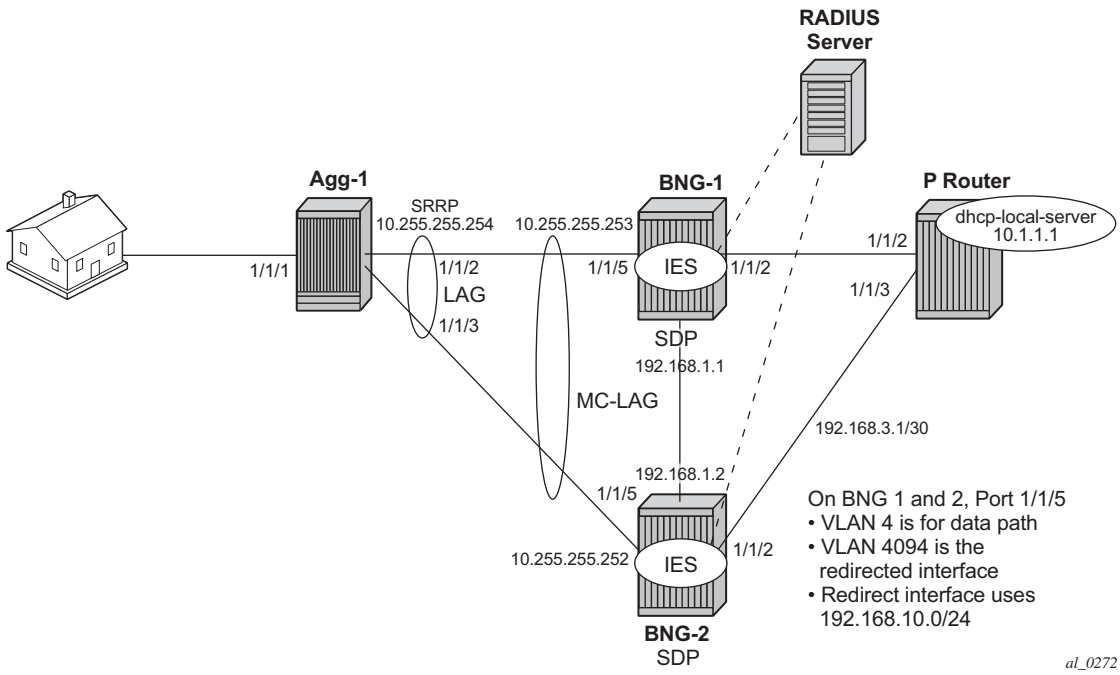

To verify the (ESM enabled) group-interface and the redirect interface are ready for multicast, use the show commands as indicated below. Remember the IES service ID is 1, the group-interface name is **group-int-1** and the interface name is **lag-redirected**.

Step 1. Verify if the group-interface and redirected interface have IGMP enabled.

```
*A:BNG-1> show router igmp group-interface
===============================================================================
IGMP Group-Interfaces
===============================================================================
FwdSvc Group-Interface                  Adm/Opr-State        Import-Policy
       SAP                              Adm/Opr-Version      Num-Groups
-------------------------------------------------------------------------------
1   group-int-1                         Up/Up                none
       lag-1:4                          3/3                  0
-------------------------------------------------------------------------------
Group-Interfaces = 1, SAPs = 1
===============================================================================
*A:BNG-1> show router igmp interface
===============================================================================
IGMP Interfaces
===============================================================================
Interface               Adm  Oper Querier         Cfg/Opr Num     Policy
                                                  Version Groups
-------------------------------------------------------------------------------
lag-redirected          Up   Down 0.0.0.0         3/3     0       none
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
```

Step 2. Ensure the subscriber is associated with an IGMP-policy. Since the IGMP-policy is associated with a subscriber-profile, verification of an IGMP-policy is performed via the sub-profile.

```
*A:BNG-1> show subscriber-mgmt sub-profile "multicast-profile-1"
===============================================================================
Subscriber Profile multicast-profile-1
===============================================================================
Description      : (Not Specified)
I. Sched. Policy : N/A
E. Sched. Policy : N/A                              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      : igmp-policy-1
Sub. MCAC Policy : N/A
NAT Policy       : N/A
Def. Encap Offset: none                             Encap Offset Mode: none
```

```
Avg Frame Size   : N/A
Preference       : 5
-------------------------------------------------------------------------------
HSMDA-2
-------------------------------------------------------------------------------
I. Qos Policy    : 1                                E. Qos Policy    : 1
                                                    E. Agg Rate Limit: Max
E. WRR Policy    : N/A                              Pkt Byte Offset  : add 0*
-------------------------------------------------------------------------------
Last Mgmt Change : 05/14/2013 10:12:49
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

After the verification, the BNGs are ready to deliver multicast streams. Next, initiate an IGMP report from a subscriber requesting a multicast channel. In this example, IGMPv3 with SSM is used. If the IPoE subscriber is receiving multicast through the subscriber SAP then the IGMP group will be associated with the SAP. Since redirection is used, the IGMP group is associated with the redirected interface instead. The output below shows that when an IGMP message is received and processed, an (S,G) binding is associated with the redirected interface. The example uses an IGMPv3 SSM message requesting (192.168.4.2, 239.255.1.1). The subscriber IP address is 10.0.0.2.

```
*A:BNG-1> show router igmp group
===============================================================================
IGMP Interface Groups
===============================================================================

(192.168.4.2,239.255.1.1)               Up Time : 0d 00:00:12
    Fwd List  : lag-redirected
===============================================================================
IGMP Host Groups
===============================================================================
(192.168.4.2,239.255.1.1)
    Fwd List  : 10.0.0.2                Up Time : 0d 00:00:12
===============================================================================
IGMP SAP Groups
===============================================================================
-------------------------------------------------------------------------------
(*,G)/(S,G) Entries : 2
===============================================================================
```

Next, verify the individual subscribers and their IGMP information. First verify the IGMP policy related to the subscriber.

```
*A:BNG-1> show service active-
subscribers igmp detail
===============================================================================
Active Subscribers Detail
===============================================================================
Subscriber                      IGMP-Policy
  HostAddr                      GrpItf                          NumGroups
    GrpAddr                     Type            Up-Time         Mode
      SrcAddr                   Type                            Blk/Fwd
```

```
-------------------------------------------------------------------------------
video_user_01                   igmp-policy-1
  10.0.0.2                      sub-int-1                 1
    239.255.1.1                 Dynamic         0d 00:01:26      Include
      192.168.4.2               Dynamic                          Fwd
-------------------------------------------------------------------------------
Number of Subscribers : 1
===============================================================================
```

Since the IGMP-policy controls bandwidth, interoperability and restricts multicast groups, it is useful to view what is defined in the IGMP-policy if the subscriber fails to receive multicast streams.

```
*A:BNG-1> show subscriber-mgmt igmp-policy "igmp-policy-1"
===============================================================================
IGMP Policy igmp-02
===============================================================================
Import Policy                       :
Admin Version                       : 3
Num Subscribers                     : 1
Host Max Group                      : No Limit
Host Max Sources                    : No Limit
Host Max Group Sources              : No Limit
Fast Leave                          : yes
Redirection Policy                  : mcast_redirect_if
Per Host Replication                : no
Egress Rate Modify                  : no
Mcast Reporting Destination Name    :
Mcast Reporting Admin State         : Disabled
===============================================================================
```

Below is a command to view the (S,G)s that all subscribers are requesting. Notice that the operational status for the host is not forwarding (notFwding), this is because multicast is not delivered directly over the subscriber SAP. All multicast traffic is delivered over the redirected interface instead.

```
*A:BNG-1> show router igmp hosts detail
===============================================================================
IGMP Host 10.0.0.2
===============================================================================
Oper Status      : notFwding  MacAddress     : 00:00:10:10:10:12
Oper version     : 3          Subscriber     : video_user_01
Num Groups       : 1          GrpItf         : sub-int-1
Max Grps Till Now: 1          IGMP-Policy    : igmp-policy-1
PPPoE SessionId  : N/A
FwdSvcId         : 1       Max Srcs Allow*: No Limit
Max Grps Allowed : No Limit   Max Grp Srcs A*: No Limit
-------------------------------------------------------------------------------
IGMP Group
-------------------------------------------------------------------------------
Group Address    : 239.255.1.1      Up Time    : 0d 00:02:38
Expires          : Not running    Mode       : Include
V1 Host Timer    : Not running    Type       : Dynamic
```

```
V2 Host Timer     : Not running    Compat Mode: IGMP Version 3
Redir.SvcId       : 1           Redir.Intf : lag-redirected
-------------------------------------------------------------
Source Address      Expires         Type          Fwd/Blk
-------------------------------------------------------------
192.168.4.2           0d 00:01:42    Dynamic       Fwd
-------------------------------------------------------------------------------
Hosts : 1
===============================================================================
```

ESM IGMP PPPoE Walkthrough

The same baseline configuration is used for PPPoE subscriber. Figure 120 shows the message flow for delivery of multicast streams to PPPoE subscribers.

Figure 120 **PPPoE Multicast Flow**

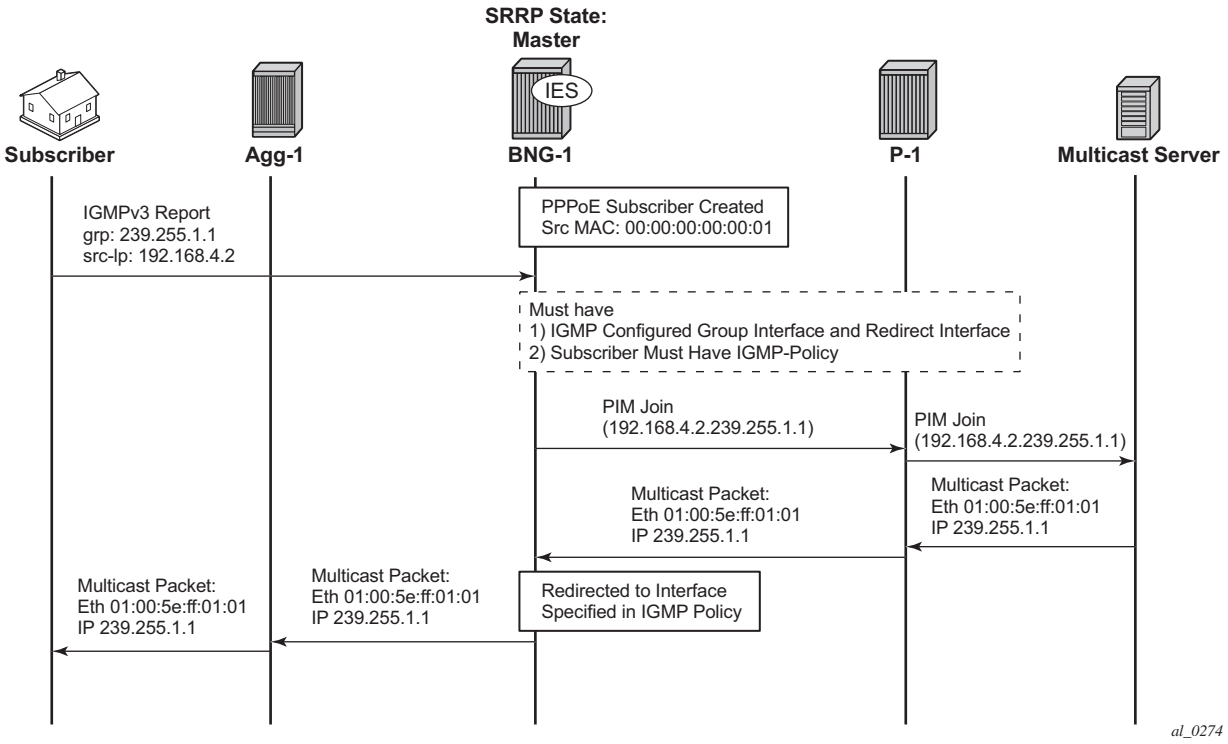

The important items are highlighted in the dotted box. By default, PPPoE subscribers receive multicast streams via Ethernet unicast over subscriber SAPs. PPPoE does not have a multicast mechanism and requires all data traffic to be unicasted. However, because multicast streams are redirected, the streams are sent as multicast at both Layers 2 and 3 (the Layer 2 header will have a multicast destination MAC address and the Layer 3 header will have a multicast destination IP address).

Verify the IGMP on the group-interface. It shows very little difference from the IPoE group interface. No multicast streams are delivered directly over the subscriber SAP group-interface.

```
*A:BNG-1> show router igmp group-interface detail
===============================================================================
IGMP Group-Interfaces
===============================================================================
FwdSvc/Grp-Intf    : 1/group-int-1
Admin-Status       : Up                 Oper-Status        : Up
Import-Policy      : none               Subnet-Check       : Enabled
Router-Alert-Check : Enabled            Sub-Hosts-Only     : Enabled
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
-------------------------------------------------------------------------------
 SAP               : lag-1:4
 Admin/Oper version: 3/3                Num Groups         : 0
 Max Groups Allowed: No Limit           Max Groups Till Now: 0
 Max Sources Allow*: No Limit
 Max Grp Srcs Allo*: No Limit
-------------------------------------------------------------------------------
Group-Interfaces = 1, SAPs = 1
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

All multicast streams should be delivered over the redirected interface. The output below shows the IGMP group for a PPPoE subscriber and also that the multicast stream is associated with the redirected interface. The (S,G) is (192.168.4.2, 239.255.1.1) and the subscriber IP address is 10.0.0.2.

```
*A:BNG-1> show router igmp group
===============================================================================
IGMP Interface Groups
===============================================================================
(192.168.4.2,239.255.1.1)                Up Time : 0d 00:05:15
    Fwd List  : lag-redirected
===============================================================================
IGMP Host Groups
===============================================================================
(192.168.4.2,239.255.1.1)
    Fwd List  : 10.0.0.2                 Up Time : 0d 00:05:15
===============================================================================
IGMP SAP Groups
===============================================================================
-------------------------------------------------------------------------------
(*,G)/(S,G) Entries : 2
===============================================================================
```

The following output shows all the subscribers and the (S,G)s they have joined. Note that there is only one PPPoE subscriber and the multicast stream is redirected.

```
*A:BNG-1> show router igmp hosts detail
===============================================================================
IGMP Host 10.0.0.2
===============================================================================
Oper Status      : Up         MacAddress     : 52:e0:50:bd:00:00
Oper version     : 3          Subscriber     : user-ppp-1
Num Groups       : 1          GrpItf         : group-int-1
Max Grps Till Now: 1          IGMP-Policy    : igmp-policy-1
PPPoE SessionId  : 1          Next query time: 0d 00:01:47
FwdSvcId         : 1          Max Srcs Allow*: No Limit
Max Grps Allowed : No Limit   Max Grp Srcs A*: No Limit
-------------------------------------------------------------------------------
IGMP Group
-------------------------------------------------------------------------------
Group Address    : 239.255.1.1    Up Time    : 0d 00:00:36
Expires          : Not running    Mode       : Include
V1 Host Timer    : Not running    Type       : Dynamic
V2 Host Timer    : Not running    Compat Mode: IGMP Version 3
Redir.SvcId      : 1              Redir.Intf : lag-redirected
-------------------------------------------------------------
Source Address     Expires        Type         Fwd/Blk
-------------------------------------------------------------
192.168.4.2        0d 00:04:03    Dynamic      Fwd
-------------------------------------------------------------------------------
Hosts : 1
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

To view the (S,G)s of a single subscriber, use the following command.

```
*A:BNG-1> show service active-subscribers igmp subscriber "user-ppp-1" detail
===============================================================================
Active Subscribers Detail
===============================================================================
Subscriber                      IGMP-Policy
  HostAddr                      GrpItf                          NumGroups
    GrpAddr                     Type            Up-Time         Mode
      SrcAddr                   Type                            Blk/Fwd
-------------------------------------------------------------------------------
user-ppp-1                    igmp-policy-1
  10.0.0.2                    group-int-1                     1
    239.255.1.1               Dynamic         0d 00:02:07     Include
      192.168.4.2             Dynamic                         Fwd
-------------------------------------------------------------------------------
Number of Subscribers : 1
===============================================================================
```

ESM IGMP MCS

The BNGs are configured with SRRP for both IPoE and PPPoE subscribers. This provides stateful redundancy when the master BNG fails. The SRRP master BNG will be the only BNG processing and answering IGMP messages, while the standby BNG synchronizes the state information of all subscribers via MCS in real time. In the event of a failure, the standby takes over and starts processing all IGMP messages. As the standby BNG has the full state information of all subscribers, including the (S,G)s they have joined, PIM starts sending joins for those (S,G)s immediately after failover. Restoration of all multicast streams happens quickly and relies on the PIM configuration and the underlying routing infrastructure. Note that the PIM command *non-dr-attract-traffic* can be used to reduce the failover outage by attracting multicast to the non designated PIM router.

The following output shows the items that are synchronized between the BNGs. To reduce the ESM multicast restoration time, it is important that all subscriber related data (IPoE, PPPoE, SRRP and IGMP) are kept in sync. BNG-1 has system IP address 192.0.2.1 and BNG-2 has system IP address 192.0.2.2.

```
*A:BNG-1> show redundancy multi-chassis sync peer 192.0.2.2 detail
===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 192.0.2.2
Description             : (Not Specified)
Authentication          : Disabled
Source IP Address       : 192.0.2.1
Admin State             : Enabled
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     : IGMP SUBMGMT-IPOE SUBMGMT-PPPOE SRRP
Sync Admin State        : Up
Sync Oper State         : Up
DB Sync State           : inSync
Num Entries             : 15
Lcl Deleted Entries     : 0
Alarm Entries           : 0
Rem Num Entries         : 15
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
===============================================================================
MCS Application Stats
===============================================================================
Application             : igmp
Num Entries             : 1
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 1
Rem Lcl Deleted Entries : 0
```

```
Rem Alarm Entries       : 0
-------------------------------------------------------------------------------
Application             : subMgmtIpoe
Num Entries             : 1
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 1
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
-------------------------------------------------------------------------------
Application             : srrp
Num Entries             : 14
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 14
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
-------------------------------------------------------------------------------
Application             : subMgmtPppoe
Num Entries             : 1
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 1
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
-------------------------------------------------------------------------------
===============================================================================
```

To check the details of the sync data across the BNGs, a tools command giving a detailed description of the IGMP information synced across MCS can be used.

```
*A:BNG-1> tools dump redundancy multi-chassis sync-database application igmp detail

If no entries are present for an application, no detail will be displayed.

FLAGS LEGEND: ld - local delete; da - delete alarm; pd - pending global delete

Peer Ip 192.0.2.2


Application IGMP
Sap-id             Client Key
 SyncTag                          DLen  Flags    timeStamp
  deleteReason code and description
-------------------------------------------------------------------------------
lag-1:4094        Host=10.0.0.2, HostGroup=239.255.1.1
 mclagdata                        20    -- -- -- 07/03/2013 15:20:49
  0x0
lag-1:4           Group=239.255.1.1
 mclagmulticast                   20    -- -- -- 07/03/2013 15:20:49
  0x0

The following totals are for:
 peer ip ALL, port/lag ALL, sync-tag ALL, application IGMP
```

```
Valid Entries:                   2
Locally Deleted Entries:         0
Locally Deleted Alarmed Entries: 0
Pending Global Delete Entries:   0
```

ESM IGMP Debug

Debug facilities allow for real-time monitoring of events happening on the system. This includes tools for debugging ESM multicast streams.

First enable the required debug on the system, then send an IGMP message to join a multicast group (S,G). The message used in this example is an IGMPv3 message with SSM.

Below is the debug information for an ESM IGMP report message at packet level.

```
debug
    router
        igmp
            packet mode egr-ingr-and-dropped
        exit
    exit

2977 2013/05/23 13:01:45.43 EST MINOR: DEBUG #2001 IGMP[9]
"IGMP[9]: RX-PKT
[012 03:58:58.090] IGMP host 10.0.0.2 V3 PDU: 10.0.0.2 -> 224.0.0.22 pduLen
20
    Type: V3 REPORT maxrespCode 0x0 checkSum 0xddf7
    Num Group Records: 1
        Group Record 0
        Type: ALW_NEW_SRCS, AuxDataLen 0, Num Sources 1
        Mcast Addr: 239.255.1.1
        Source Address List
            192.168.4.2

"
```

Below is the debug information for an ESM IGMP host. Notice the multicast stream is redirected to the LAG interface and that an MCS entry is installed for the new IGMP group.

```
debug
    router
        igmp
            host "10.0.0.2"
        exit
    exit


9 2013/07/03 15:26:32.74 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9]
"IGMP[ies1 inst 9]: igmpIfGroupAdd
```

```
Adding 239.255.1.1 to IGMP host 10.0.0.2 database"

10 2013/07/03 15:26:32.74 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9]
"IGMP[ies1 inst 9]: igmpProcessGroupRec
Process group rec ALW_NEW_SRCS received on host 10.0.0.2 for group 239.255.1.1 i
n mode INCLUDE. Num srcs 1"

11 2013/07/03 15:26:32.74 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9]
"IGMP[ies1 inst 9]: igmpIfSrcAdd
Adding i/f source entry for host 10.0.0.2 (192.168.4.2,239.255.1.1) to IGMP fwdList
Database, redir if interface lag-redirected [ifIndex 16]"

12 2013/07/03 15:26:32.73 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9]
"IGMP[ies1 inst 9]: igmpMcsAddIfGroup
Building MCS entry for host 10.0.0.2, group 239.255.1.1"
```

Below is the debug information for ESM IGMP when MCS sync is enabled. The MCS sends a sync message for the redirect interface.

```
debug
    router
        igmp
            mcs "lag-redirected"
        exit
    exit


20 2013/07/03 15:28:26.20 EST MINOR: DEBUG #2001 ies1 IGMP MCS[9]
"IGMP MCS[9]: TX-MCS Data
interface lag-redirected [ifIndex 16]
Key Type: Group, Len: 9, Grp Addr: 239.255.1.1
Data Type: Group, Len: 20, Ver: 0, RecType: 1, Compat Mode: 3,
Num Fwd Srcs: 1, Num Blk Srcs: 0
Fwd Sources:
    192.168.4.2
"

21 2013/07/03 15:28:26.20 EST MINOR: DEBUG #2001 ies1 IGMP MCS[9]
"IGMP MCS[9]: TX-MCS Data
interface lag-redirected [ifIndex 16]
Key Type: Group, Len: 9, Grp Addr: 239.255.1.1
Data Type: Group, Len: 20, Ver: 0, RecType: 1, Compat Mode: 3,
Num Fwd Srcs: 1, Num Blk Srcs: 0
Fwd Sources:
    192.168.4.2
"
```

The corresponding debug information for ESM IGMP MCS sync on BNG-2 looks as follows:

```
2 2013/07/03 20:30:24.97 UTC MINOR: DEBUG #2001 ies1 IGMP MCS[5]
"IGMP MCS[5]: RX-MCS Data
interface lag-redirected [ifIndex 15]
Key Type: Group, Len: 9, Grp Addr: 239.255.1.1
```

```
Data Type: Group, Len: 20, Ver: 0, RecType: 1, Compat Mode: 3,
Num Fwd Srcs: 1, Num Blk Srcs: 0
Fwd Sources:
    192.168.4.2
"
```

The same debug commands can be used for viewing IGMP leave messages. Below is the debug information for an ESM IGMP leave at the packet level. The leave report message received over the subscriber SAP results in the multicast stream being stopped on the redirected interface, after ensuring no other CPE devices still require the multicast streams (by means of a query).

```
debug
    router
        igmp
            packet mode egr-ingr-and-dropped
        exit
    exit


37 2013/07/03 15:32:10.05 EST MINOR: DEBUG #2001 ies1 IGMP[9]
"IGMP[9]: RX-PKT
[001 03:23:17.050] IGMP host 10.0.0.2 V3 PDU: 10.0.0.2 -> 224.0.0.22 pduLen
20
    Type: V3 REPORT maxrespCode 0x0 checkSum 0xddf3
    Num Group Records: 1
        Group Record 0
        Type: BLK_OLD_SRCS, AuxDataLen 0, Num Sources 1
        Mcast Addr: 239.255.1.1
        Source Address List
            192.168.4.2

"

38 2013/07/03 15:32:10.05 EST MINOR: DEBUG #2001 ies1 IGMP[9]
"IGMP[9]: TX-PKT
[001 03:23:17.050] IGMP interface lag-
redirected [ifIndex 16] V3 PDU: 192.168.10.253
-> 239.255.1.1 pduLen 16
    Type: QUERY maxrespCode 0xa checkSum 0xf26d
    GroupAddr: 239.255.1.1
        S bit 0, QRV 2, Encoded-QQIC 125, NumSources 1
        Source Address List:
            192.168.4.2
"

39 2013/07/03 15:32:11.36 EST MINOR: DEBUG #2001 ies1 IGMP[9]
"IGMP[9]: TX-PKT
[001 03:23:18.370] IGMP interface lag-
redirected [ifIndex 16] V3 PDU: 192.168.10.253
-> 239.255.1.1 pduLen 16
    Type: QUERY maxrespCode 0xa checkSum 0xf26d
    GroupAddr: 239.255.1.1
        S bit 0, QRV 2, Encoded-QQIC 125, NumSources 1
        Source Address List:
            192.168.4.2
"
```

Below is the debug information for an ESM IGMP host showing various IGMP events. The MCS also signals the removal of the IGMP entry in the database.

```
debug
    router
        igmp
            host "192.168.0.10"
        exit
    exit

44 2013/07/03 15:33:06.00 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9]
"IGMP[ies1 inst 9]: igmpProcessGroupRec
Process group rec BLK_OLD_SRCS received on host 10.0.0.2 for group 239.255.1.1 i
n mode INCLUDE. Num srcs 1"

45 2013/07/03 15:33:06.00 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9]
"IGMP[ies1 inst 9]: igmpProcessIfSrcTimerExp
Source Timer expired for IGMP host 10.0.0.2 (192.168.4.2,239.255.1.1)"

46 2013/07/03 15:33:06.00 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9]
"IGMP[ies1 inst 9]: igmpIfSrcDel
Deleting i/f source entry for host 10.0.0.2 (192.168.4.2,239.255.1.1) from IGMP Data
base. DeleteFromAvl: 1 Redir 0"

47 2013/07/03 15:33:06.00 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9]
"IGMP[ies1 inst 9]: igmpIfGroupDel
Deleting 239.255.1.1 from IGMP host 10.0.0.2 database"

48 2013/07/03 15:33:05.99 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9]
"IGMP[ies1 inst 9]: igmpMcsDelIfGroup
Deleting MCS entry for host 10.0.0.2, group 239.255.1.1, Glb"

49 2013/07/03 15:33:05.99 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9]
"IGMP[ies1 inst 9]: igmpMcsDelIfGroup
Deleting MCS entry for host 10.0.0.2, group 239.255.1.1, Glb"

50 2013/07/03 15:33:06.00 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9]
"IGMP[ies1 inst 9]: igmpMcsDelIfGroup
Deleting MCS entry for host 10.0.0.2, group 239.255.1.1, Glb"
```

The debug information when MCS removes the entry on BNG-1 is shown below. Notice MCS also triggers the backup BNG to remove the multicast stream.

```
debug
    router
        igmp
            mcs "group-int-1"
        exit
    exit

69 2013/07/03 15:34:42.43 EST MINOR: DEBUG #2001 ies1 IGMP MCS[9]
"IGMP MCS[9]: TX-MCS Data
interface lag-redirected [ifIndex 16]
Key Type: Group, Len: 9, Grp Addr: 239.255.1.1
Data Type: Group, Len: 20, Ver: 0, RecType: 1, Compat Mode: 3,
```

```
Num Fwd Srcs: 1, Num Blk Srcs: 0
Fwd Sources:
    192.168.4.2
"

70 2013/07/03 15:34:42.43 EST MINOR: DEBUG #2001 ies1 IGMP MCS[9]
"IGMP MCS[9]: TX-MCS Data
interface lag-redirected [ifIndex 16]
Key Type: Group, Len: 9, Grp Addr: 239.255.1.1
Data Type: Group, Len: 20, Ver: 0, RecType: 1, Compat Mode: 3,
Num Fwd Srcs: 1, Num Blk Srcs: 0
Fwd Sources:
    192.168.4.2
"

71 2013/07/03 15:34:44.36 EST MINOR: DEBUG #2001 ies1 IGMP MCS[9]
"IGMP MCS[9]: TX-MCS Data (GlblDel)
interface lag-redirected [ifIndex 16]
Key Type: Group, Len: 9, Grp Addr: 239.255.1.1
Data Type: Group, Len: 16, Ver: 0, RecType: 1, Compat Mode: 3,
Num Fwd Srcs: 0, Num Blk Srcs: 0
"

72 2013/07/03 15:34:44.37 EST MINOR: DEBUG #2001 ies1 IGMP MCS[9]
"IGMP MCS[9]: TX-MCS Data (GlblDel)
interface lag-redirected [ifIndex 16]
Key Type: Group, Len: 9, Grp Addr: 239.255.1.1
Data Type: Group, Len: 16, Ver: 0, RecType: 1, Compat Mode: 3,
Num Fwd Srcs: 0, Num Blk Srcs: 0
"
```

The debug information on BNG-2 shows the sync message received over MCS for the removal of the multicast (S,G).

```
13 2013/07/03 20:34:44.37 UTC MINOR: DEBUG #2001 ies1 IGMP MCS[5]
"IGMP MCS[5]: RX-MCS Data
interface lag-redirected [ifIndex 15]
Key Type: Group, Len: 9, Grp Addr: 239.255.1.1
Data Type: Group, Len: 20, Ver: 0, RecType: 1, Compat Mode: 3,
Num Fwd Srcs: 1, Num Blk Srcs: 0
Fwd Sources:
    192.168.4.2
"
```

Conclusion

Multicast is an essential part of Triple Play Services. The SR 7750 TPSDA solution is much more than a baseline multicast delivery, it includes individual subscriber awareness and offers a full state redundancy option. Subscriber awareness allows for fine tuning of subscriber multicast settings and for troubleshooting on a per subscriber basis. Full state redundancy reduces failover time and ensures high availability of multicast services. This example provided a complete configuration walk through of both the IPoE and PPPoE SRRP model with redirection. All multicast streams can be redirected to a dedicated interface for all subscribers to receive.

ESM IPv4: Multicast with SRRP

This chapter describes ESM IPv4 multicast with SRRP configurations.

Topics in this chapter include:

Applicability

This example is applicable to all 7750 SR-12 with IOM3-XP and IMMs, and needs chassis mode c as a minimum. This is also supported on 7450 ESS chassis in mixed-mode and 7750 SR-c4/12 platform.

The configuration was tested on release 11.0R1 and covers both IPoE and PPPoE subscribers.

Overview

Triple Play Service Delivery Architecture (TPSDA) has allowed operators to integrate High Speed Internet (HSI), voice, and video services within a single network infrastructure. The goal of this chapter is to walk through the configuration of a redundant TPSDA multicast architecture and the configuration of advanced multicast filters. The topics are divided into the following sections:

- Enhanced Subscriber Management (ESM) multicast baseline configuration
  - IGMP configuration on ESM group interface
  - ESM IGMP-policy configuration
- PPPoE ESM multicast configuration
- IPoE ESM multicast configuration
- IGMP Subscriber Router Redundancy Protocol (SRRP)
  - Multi-Chassis Synchronization (MCS) walkthrough
- Advanced ESM IGMP configurations

– Filter list

The network topology displayed in Figure 121 shows a typical TPSDA setup. It consists of three 7750s and a single 7450. Two 7750s are configured as Broadband Network Gateways (BNGs) and the third 7750 is configured as a **P** router. The 7450 is used as an aggregation switch to aggregate all subscribers.

Figure 121 **Network Topology Overview**

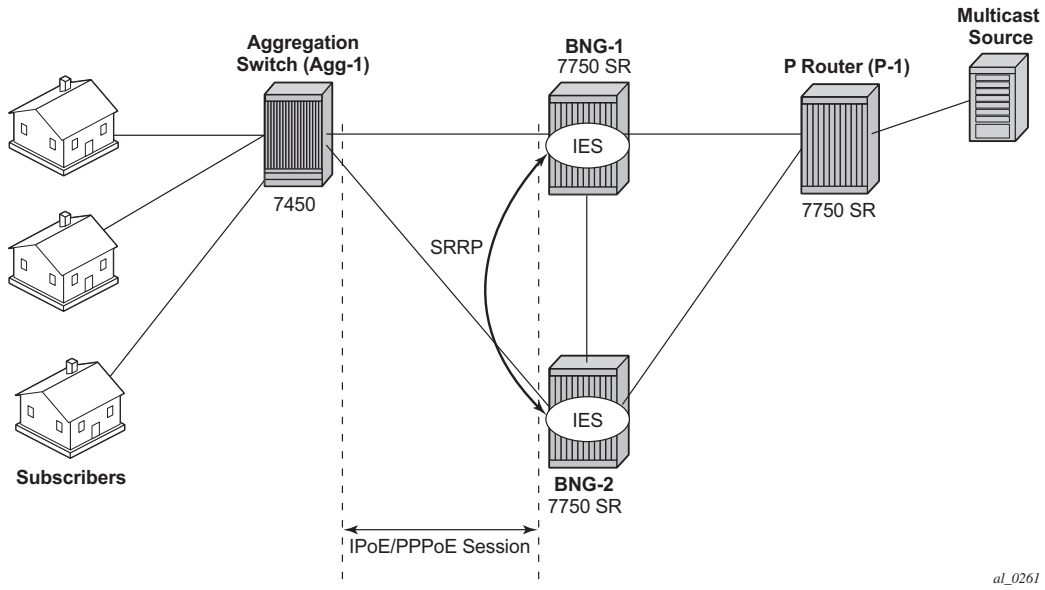

Both BNGs are configured with SRRP to provide redundancy. Note that SRRP is only used for redundancy purposes. SRRP is not mandatory for supporting multicast. The P router is connected to the multicast source and to the network side of both BNGs. The connections between the BNGs and the P router are also running PIM to provide multicast delivery. On the access side, the two BNGs are connected to an aggregation switch which aggregates the traffic originating from both PPPoE and IPoE subscribers. The BNGs are IGMP capable and will respond to subscribers' IGMP requests.

There are two requirements to enable multicast delivery using ESM. First, the ESM group interface must have IGMP enabled. Second, the ESM subscribers must be configured with an IGMP-policy to receive multicast. When both requirements are met, the BNG will process the subscribers' IGMP messages, otherwise, IGMP messages are simply ignored and dropped. All customer premise device (CPE) IGMP messages are aggregated via the 7450 and passed to the BNGs. Since the

BNGs are running SRRP, the SRRP master is the only BNG processing and answering the IGMP messages. Protocol Independent Multicasting (PIM) is then used between the BNG and the P router to request the multicast content. If PIM is successful in retrieving the multicast group, the multicast stream is forwarded towards the subscribers. This is the typical multicast delivery model for TPSDA.

Configuration

This section expects a basic knowledge of ESM.

ESM SRRP Configuration Overview

Figure 122 shows the addressing scheme used in the setup. The example uses numbered SRRP subscriber interfaces with static SAPs serving both IPoE and PPPoE subscribers. The configuration of the RADIUS server is out of the scope of this example.

Figure 122 **Network Topology Used for the Testing**

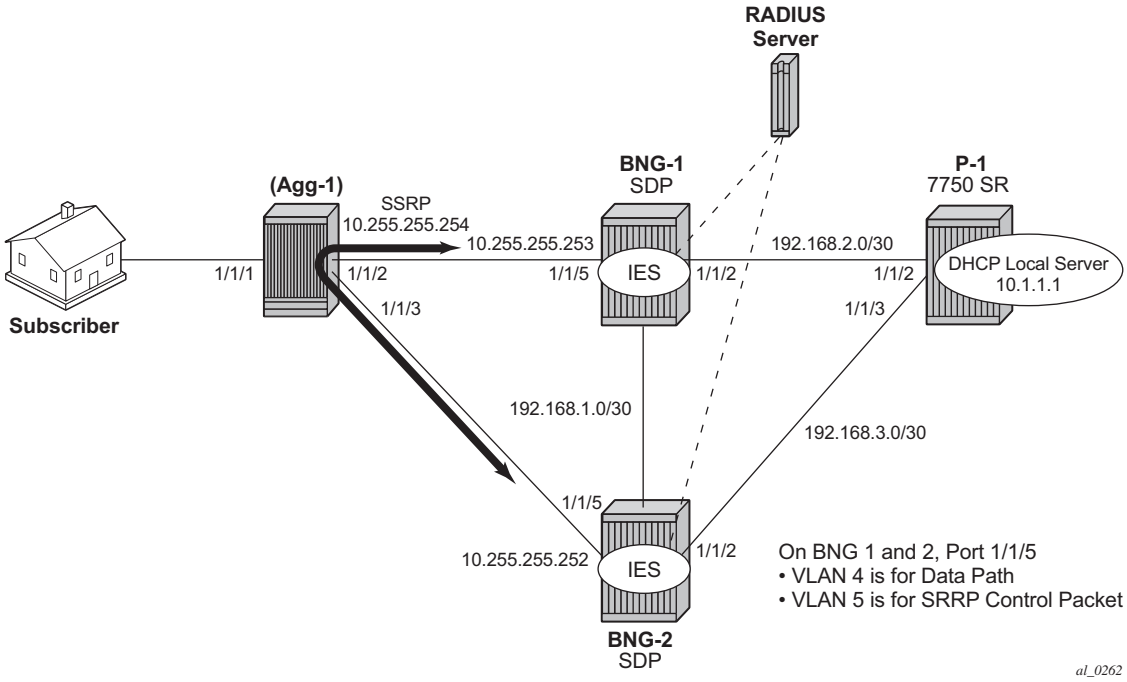

The baseline configuration for BNG-1 is shown below without any IGMP configuration. The subscriber-interface is configured in an IES, though it is also possible to configure the subscriber-interface in a VPRN. OSPF and PIM are also provisioned to provide routing and multicast capabilities. The SRRP configuration with priority 100 ensures BNG-1 is the master when both BNGs are active as the SRRP priority for BNG-2 is lower.

```
*A:BNG-1>config>service>ies# info
----------------------------------------------
            description "BNG-1"
            redundant-interface "MClink-BNG-1-BNG-2" create
                address 192.168.1.0/31
                ip-mtu 1500
                spoke-sdp 1:1 create
                    no shutdown
                exit
            exit
            interface "Int-BNG-1-P-1" create
                address 192.168.2.1/30
                sap 1/1/2 create
                    no shutdown
                exit
            exit
            subscriber-interface "sub-int-1" create
                address 10.255.255.253/8 gw-ip-address 10.255.255.254 track-srrp 1
                group-interface "group-int-1" create
                    dhcp
                        server 192.168.0.1
                        lease-populate 10
                        client-applications dhcp ppp
                        gi-address 10.255.255.253
                        no shutdown
                    exit
                    authentication-policy "auth-policy-1"
                    redundant-interface "MClink-BNG-1-BNG-2"
                    sap 1/1/5:4 create
                        sub-sla-mgmt
                            def-sub-id use-sap-id
                            def-sub-profile "multicast-profile-1"
                            def-sla-profile "sla-profile-1"
                            sub-ident-policy "sub-ident-policy-1"
                            multi-sub-sap 10
                            no shutdown
                        exit
                    exit
                    sap 1/1/5:5 create
                    exit
                    srrp 1 create
                        message-path 1/1/5:5
                        priority 100
                        no shutdown
                    exit
                    pppoe
                        no shutdown
                    exit
                exit
            exit
```

```
A:BNG-1>config>router# info
        ospf
            traffic-engineering
            area 0.0.0.0
                interface "system"
                    no shutdown
                exit
                interface "int-BNG-1-BNG-2"
                    interface-type point-to-point
                    metric 10000
                    no shutdown
                exit
                interface "sub-int-1"
                    no shutdown
                exit
                interface "int-BNG-1-P-1"
                    no shutdown
                exit
            exit
        exit
        pim
            interface "nt-BNG-1-P-1
            exit
            no shutdown
```

The baseline configuration for BNG-2 is shown below without any IGMP configuration. The default SRRP priority for BNG-2 is lower than the SRRP priority for BNG-1 and hence BNG-2 will be in standby mode.

```
*A:BNG-2>config>service>ies# info
----------------------------------------------
            description "BNG-2"
            redundant-interface "MClink-BNG-2-BNG-1" create
                address 192.168.1.1/31
                ip-mtu 1500
                spoke-sdp 1:1 create
                    no shutdown
                exit
            exit
            interface "int-BNG-2-P-1" create
                address 192.168.3.1/30
                sap 1/1/2 create
                    no shutdown
                exit
            exit
            subscriber-interface "sub-int-1" create
                address 10.255.255.252/8 gw-ip-address 10.255.255.254 track-srrp 1
                group-interface "group-int-1" create
                    dhcp
                        server 192.168.0.1
                        lease-populate 10
                        client-applications dhcp ppp
                        gi-address 10.255.255.252
                        no shutdown
                    exit
```

```
                        authentication-policy "auth-policy-1"
                        redundant-interface "MClink-BNG-2-BNG-1"
                        sap 1/1/5:4 create
                            sub-sla-mgmt
                                def-sub-id use-sap-id
                                def-sub-profile "multicast-profile-1"
                                def-sla-profile "sla-profile-1"
                                sub-ident-policy "sub-ident-policy-1"
                                multi-sub-sap 10
                                no shutdown
                            exit
                        exit
                        sap 1/1/5:5 create
                        exit
                        srrp 1 create
                            message-path 1/1/5:5
                            no shutdown
                        exit
                        pppoe
                            no shutdown
                        exit

                    exit
                exit

*A:BNG-2>config>router# info
            ospf
                traffic-engineering
                area 0.0.0.0
                    interface "system"
                        no shutdown
                    exit
                    interface "int-BNG-2-BNG-1"
                        interface-type point-to-point
                        metric 10000
                        no shutdown
                    exit
                    interface "sub-int-1"
                        no shutdown
                    exit
                    interface "int-BNG-2-P-1"
                        no shutdown
                    exit
                exit
            exit
            pim
                interface "int-BNG-2-P-1"
                exit
                no shutdown
            exit
```

The baseline configuration for the 7450 aggregation switch is shown below. Two VPLS services are configured. The first VPLS, VPLS 1, is responsible for passing the SRRP control traffic over VLAN 5. The second VPLS, VPLS 2, is responsible for passing all subscriber data traffic over VLAN 4.

```
*A:Agg-1>config>service>info
            vpls 1 customer 1 create
                sap 1/1/2:5 create
                    no shutdown
                exit
                sap 1/1/3:5 create
                    no shutdown
                exit
            no shutdown
            exit
            vpls 2 customer 1 create
                sap 1/1/2:4 create
                    no shutdown
                exit
                sap 1/1/3:4 create
                    no shutdown
                exit
                sap 1/1/1:4 create
                    no shutdown
                exit
            no shutdown
            exit
```

The baseline configuration on the P router is shown below. The P router has a local DHCP server configured and performs the DHCP address assignment. It is also attached to the multicast source and uses PIM to deliver multicast streams.

```
*A:P-1>config>router>info
#------------------------------------------------
echo "Local DHCP Server Configuration"
#------------------------------------------------
        dhcp
            local-dhcp-server "dhcp-local-server" create
                use-gi-address scope pool
                pool "pool-1" create
                    subnet 10.0.0.0/8 create
                        options
                            subnet-mask 255.0.0.0
                            default-router 10.255.255.254
                        exit
                        address-range 10.0.0.10 10.0.0.254
                    exit
                exit
                no shutdown
            exit
        exit
#------------------------------------------------
echo "IP Configuration"
#------------------------------------------------
        interface "dhcp-lb1"
            address 192.168.0.1/32
            loopback
            local-dhcp-server "dhcp-local-server
            no shutdown
        exit
        interface "int-P-1-BNG-1"
            address 192.168.2.2/30
```

```
            port 1/1/2
            no shutdown
        exit
        interface "int-P-1-BNG-2"
            address 192.168.3.2/30
            port 1/1/3
            no shutdown
        exit
        interface "P-1-multicast-source"
            address 192.168.4.1/30
            port 1/1/1
            no shutdown
        exit
        interface "system"
            address 192.0.2.3/32
            no shutdown
        exit

#--------------------------------------------------
        ospf
            area 0.0.0.0
                interface "system"
                    no shutdown
                exit
                interface "int-P-1-BNG-1"
                    no shutdown
                exit
                interface "int-P-1-BNG-2"
                    no shutdown
                exit
                interface "P-1-multicast-source"
                    no shutdown
                exit
            exit
        exit
        pim
            interface "int-P-1-BNG-1"
            exit
            interface "int-P-1-BNG-2"
            exit
            interface "P-1-multicast-source"
            exit
        exit
```

Enable IGMP on Group Interfaces

The configuration below adds the group interface to IGMP. If the subscriber-interface is configured in a VPRN, each VPRN will have its individual IGMP instance. Add the group-interface to the IGMP instance.

```
*A:BNG-1>config>router>igmp# info
----------------------------------------------
               group-interface "group-int-1"
                    no shutdown
```

```
                exit
```

Placing the group-interface into IGMP is the first step required to deliver multicast content. The options available in this IGMP context can be classified into two categories:

1. Bandwidth and multicast group management
2. Interoperability

```
[no] disable-router* - Enable/disable the IGMP router alert check option
[no] import          - Import a policy to filter IGMP packets
[no] max-groups      - Configure the maximum number of groups for this
                       group-interface
[no] max-grp-sources - Configure the maximum number of group sources for this
                       group-interface
[no] max-sources     - Configure the maximum number of sources for this
                       group-interface
[no] mcac            + Configure multicast CAC policy and constraints for this
                       interface
[no] shutdown        - Administratively enable/disable the interface
[no] sub-hosts-only  - Enable/disable the IGMP traffic from known hosts only
[no] subnet-check    - Enable/disable local subnet checking for IGMP
[no] version         - Configure the version of IGMP
```

The bandwidth and multicast group management options are:

- Import — Used for white-listing or black-listing multicast groups in the IGMP control plane. More configuration detail is offered in a later section.
- Max-groups — Controls the maximum number of groups (channels) allowed on the group interface.
- Max-sources — Controls the maximum number of sources of the multicast streams on a group interface.
- Max-grp-sources — Specifies the maximum number of multicast group and source pairs for a group-interface.
- MCAC — Multicast Connection Admission Control (MCAC) is a bandwidth management feature to control the amount of multicast content a group interface is allowed to receive. It can also be applied at subscriber level to offer a hierarchical control. Multicast bandwidth can be controlled on a per subscriber basis.

The interoperability options available are:

- Router alert — enable or disable router alert processing.

- Sub-hosts-only — Only subscriber originated IGMP messages are accepted and anything else is ignored. Sometimes, IGMP message might not arrive directly from the subscriber. For example, an aggregation switch or DSLAM residing between the CPE and the BNG might perform IGMP proxy. The switch/DSLAM will insert its own source IP-address in place of the subscriber.

  It should be noted that, when an IGMP proxy is used, the identity of the subscriber is lost (since the original source IP of the IGMP message is replaced).
- Subnet-check — IGMP messages will be checked against the group interface subnet. All IGMP messages with a source address that is not in the local subnet are dropped.
- Version — The RFCs define 3 versions of IGMP, all of them are supported by SROS.

  It must be noted that when subscribers are sending IGMPv1 or v2 in a bridged LAN, suppression of IGMP messages can occur. If an IGMP host detects the presence of another host reporting for the same multicast group, it will suppress its own IGMP report message and silently receive the multicast stream. When IGMP messages are suppressed, the BNG might not be able to account for the real multicast bandwidth consumption of each subscriber. IGMPv3, on the other hand, forces all hosts to send IGMP reports. This guarantees that the BNG identifies each subscriber's IGMP request.

ESM IGMP Policy

In addition to enabling IGMP on the group interface, the subscriber must be allowed to receive multicast content through an IGMP policy. For this purpose, the IGMP-policy is associated with the subscriber's subscriber-profile. Therefore during authentication, either RADIUS, the local user database (LUDB), or the default-sub-profile should return a sub-profile with an IGMP policy. The provisioning requires two steps:

Step 1. Create the IGMP policy.

```
*A:BNG-1> config subscr-mgmt
     igmp-policy "igmp-policy-1" create
            exit
```

Step 2. Add the IGMP policy to a subscriber-profile.

```
*A:BNG-1> config subscr-mgmt
     sub-prof "multicast-profile-1"
            igmp-policy "igmp-policy-1"
```

The above configuration is the minimum requirement for a subscriber to receive multicast streams. The different options inside an IGMP policy are:

```
[no] description     - Description for the IGMP policy
[no] egress-rate-mo* - Configure the egress rate modification
[no] fast-leave      - Enable/disable IGMP fast-leave processing
[no] import          - Specify the import policy to filter IGMP packets
[no] max-num-groups  - Configure the max number of multicast groups
[no] max-num-grp-so* - Configure the max number of multicast group sources
[no] max-num-sources - Configure the max number of multicast sources
[no] mcast-reporting + Configure the mcast reporting
[no] per-host-repli* - Enable/disable IGMP per-host-replication processing
[no] redirection-po* - Specify the IGMP redirection policy
     static          + Add/remove IGMP static group membership
[no] version         - Configure the version of IGMP
```

Again, two groups of options are available: the bandwidth and multicast group management options, and the interoperability options.

Bandwidth and multicast group management options:

- Import — Used for white-listing and black-listing multicast groups. This allows the import policy to be defined per subscriber.
- Max-num-group — Limits the maximum multicast groups for the group interface. This limits the groups per subscriber.
- Max-num-sources — Limits the maximum multicast sources for the group interface. This limits the sources per subscriber.
- Max-num-grp-sources — Limits the maximum multicast group and source pairs for the group interface.
- Egress-rate-modify — This feature adjusts the subscriber queue bandwidth according to multicast consumption. It is used in conjunction with MCAC. An MCAC policy defines the bandwidth consumption per multicast group. As a subscriber joins a multicast group, the bandwidth of the multicast channel is subtracted from the subscriber queue bandwidth. The remaining bandwidth is what the subscriber can use for all other services.

Interoperability options:

- Fast-leave — Enables the router to withdraw the multicast group quickly when receiving an IGMP leave message without any last query. This should be used in a subscriber per SAP (dot1q or qinq) model.
- Static — This allows the provisioning of static multicast groups that the subscriber will receive regardless of any IGMP report. The static multicast group can be Source Specific Multicast (SSM) based.

- Per-host-replication — SR OS has the capability to replicate a multicast source per subscriber. For example, if 10 subscribers are requesting the same multicast group, then 10 multicast streams are replicated and delivered individually to each subscriber. PPPoE requires service delivery to be point to point. To achieve this, the Ethernet header destination address is the subscriber's source MAC. The IP layer destination remains as the multicast group that the subscriber requested. For IPoE, without per-host replication, the standard multicast MAC representing the multicast group is used as the destination MAC address when delivering the multicast content to the subscribers. When per-host replication is used for IPoE subscribers, the destination MAC address will be the host's own MAC address and the IP destination will remain as the multicast group that the subscriber is interested in. The multicast content is then delivered to the subscriber via MAC learning as a unicast stream.

  It should be noted that when per-host-replication is enabled, all multicast content will be using the subscriber queues. It is no longer necessary to use egress-rate-modify as mentioned above.
- Redirection-policy — Another popular model for multicast deployment is to redirect all multicast content to another interface instead of sending the content directly to the subscriber. All subscribers use a common VLAN to receive the multicast content.

ESM IGMP IPoE Walkthrough

With the baseline configuration applied, the BNG is ready to process IGMP messages and deliver multicast. Figure 123 shows a flow for IPoE subscribers requesting and receiving multicast traffic. The key items are highlighted in dotted box:

1. The ESM group-interface must have IGMP enabled
2. The subscriber must be associated with an IGMP-policy via sub-profile.

The subscriber sends an IGMPv3 report using (192.168.4.2, 239.255.1.1).

Figure 123 **IPoE Subscriber Multicast Flow**

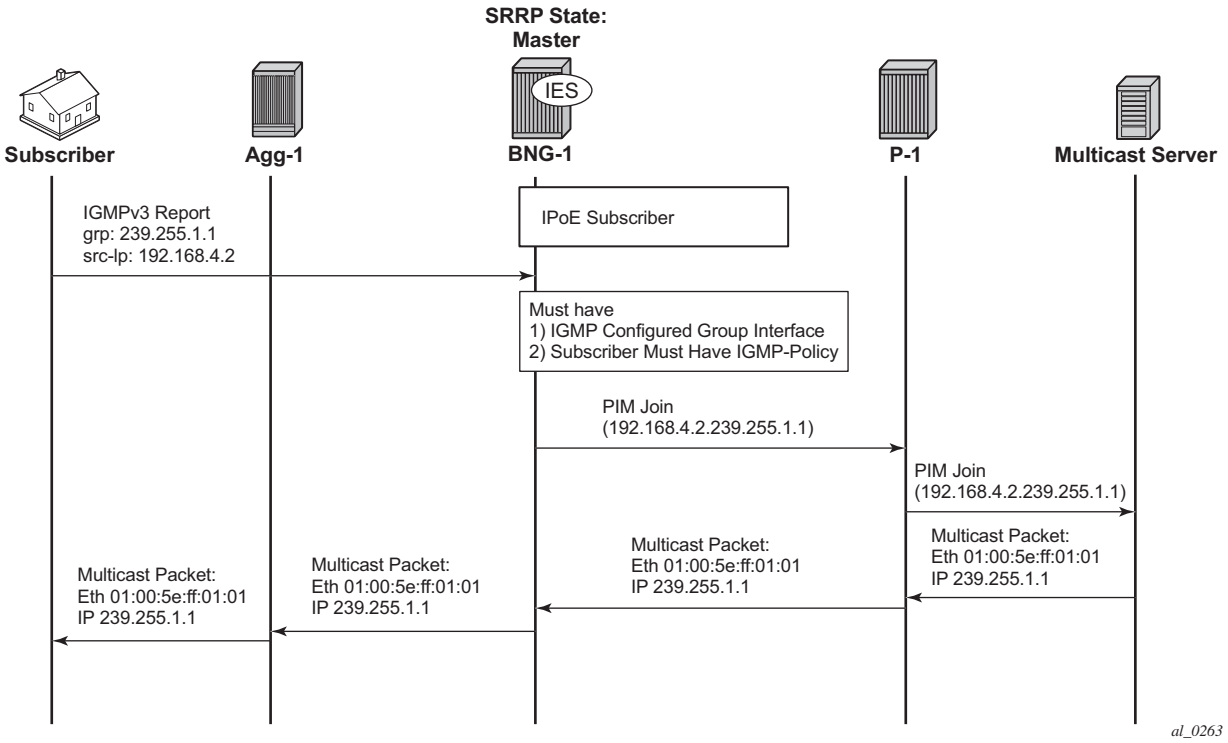

To verify that the group interface is ready for multicast, use the **show** command as indicated below. Remember that the IES service id is 1 and the group-interface name is *group-int-1*.

Step 1. Verify if the group-interface has IGMP enabled.

```
*A:BNG-1> show router igmp group-interface
===============================================================================
IGMP Group-Interfaces
===============================================================================
FwdSvc Group-Interface                 Adm/Opr-State        Import-Policy
       SAP                             Adm/Opr-Version      Num-Groups
-------------------------------------------------------------------------------
1      group-int-1                     Up/Up                none
       1/1/5:4                         3/3                  0
-------------------------------------------------------------------------------
Group-Interfaces = 1, SAPs = 1
===============================================================================
```

Step 2. Ensure the subscriber is associated with an IGMP-policy. Since the IGMP-policy is associated with a subscriber-profile, verifying the IGMP-policy is achieved via sub-profile.

```
*A:BNG-1> show subscriber-mgmt sub-profile "multicast-profile-1"
===============================================================================
Subscriber Profile multicast-profile-1
===============================================================================
Description      : (Not Specified)
I. Sched. Policy : N/A
E. Sched. Policy : N/A                              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      : igmp-policy-1
Sub. MCAC Policy : N/A
NAT Policy       : N/A
Def. Encap Offset: none                             Encap Offset Mode: none
Avg Frame Size   : N/A
Preference       : 5
-------------------------------------------------------------------------------
HSMDA-2
-------------------------------------------------------------------------------
I. Qos Policy    : 1                                E. Qos Policy    : 1
                                                    E. Agg Rate Limit: Max
E. WRR Policy    : N/A                              Pkt Byte Offset  : add 0*
-------------------------------------------------------------------------------
Last Mgmt Change : 05/14/2013 10:12:49
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

After the verification, the BNGs are ready to deliver multicast content.

First, initiate an IGMP report from a subscriber requesting a multicast channel. In this example, IGMPv3 SSM is used. IPoE by default replicates per-SAP. If the IGMP message was successfully received and processed, an (S,G) binding will be associated with the subscriber SAP.

In this case, the IGMPv3 SSM message requests (192.168.4.2, 239.255.1.1). The subscriber host is assigned an IP address of 10.0.0.24. To verify the IGMP message was successfully processed, check that the (S,G) is learned on the IGMP instance. The example below shows a successful IGMP message processed by the BNG, the (S,G) is registered against the subscriber SAP.

```
*A:BNG-1> show router igmp group
===============================================================================
IGMP Interface Groups
===============================================================================
===============================================================================
IGMP Host Groups
===============================================================================
===============================================================================
IGMP SAP Groups
```

```
===============================================================================
(192.168.4.2,239.255.1.1)
    Fwd List  : 1/1/5:4                  Up Time : 0d 00:00:08
-------------------------------------------------------------------------------
(*,G)/(S,G) Entries : 1
===============================================================================
```

For more IGMP details on the group interface, including maximum multicast groups and bandwidth management, use the following command:

```
*A:BNG-1> show router igmp group-interface detail
===============================================================================
IGMP Group-Interfaces
===============================================================================
FwdSvc/Grp-Intf    : 1/group-int-1
Admin-Status       : Up                 Oper-Status        : Up
Import-Policy      : none               Subnet-Check       : Disabled
Router-Alert-Check : Enabled            Sub-Hosts-Only     : Disabled
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
-------------------------------------------------------------------------------
 SAP               : 1/1/5:4
 Admin/Oper version: 3/3                Num Groups         : 1
 Max Groups Allowed: No Limit           Max Groups Till Now: 1
 Max Sources Allow*: No Limit
 Max Grp Srcs Allo*: No Limit
-------------------------------------------------------------------------------
  Group-Address    : 239.255.1.1          Up Time            : 0d 00:04:05
  Expires          : N/A                Mode               : include
  V1 Host Timer    : Not running        Type               : dynamic
  V2 Host Timer    : Not running        Compat Mode        : IGMP Version 3
   ---------------------------------------------------------------
   GrpSrc-Address    Expires            Type                 Fwd/Blk
   ---------------------------------------------------------------
   192.168.4.2       0d 00:03:50        dynamic              Fwd
-------------------------------------------------------------------------------
Group-Interfaces = 1, SAPs = 1
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

If the subscriber fails to receive multicast traffic, check if the subscriber has an associated IGMP policy with the following command. If the subscriber entry is missing, make sure the subscriber has a sub-profile that is tied to an IGMP-policy.

```
*A:BNG-1> show service active-
subscribers igmp detail
===============================================================================
Active Subscribers Detail
===============================================================================
Subscriber                      IGMP-Policy
  HostAddr                      GrpItf                          NumGroups
    GrpAddr                     Type            Up-Time         Mode
```

```
      SrcAddr                   Type                                Blk/Fwd
-------------------------------------------------------------------------------
Subscriber-1                    igmp-policy-1
-------------------------------------------------------------------------------
Number of Subscribers : 1
```

Another possibility for failing to receive multicast traffic could be due to the control mechanisms inside the IGMP-policy such as: bandwidth control, multicast groups restrictions, and interoperability options. Use the following command to view the IGMP policy configured control parameters.

```
*A:BNG-1> show subscriber-mgmt igmp-policy "igmp-policy-1"
===============================================================================
IGMP Policy igmp-policy-1
===============================================================================
Import Policy                       :
Admin Version                       : 3
Num Subscribers                     : 0
Host Max Group                      : No Limit
Host Max Sources                    : No Limit
Host Max Group Sources              : No Limit
Fast Leave                          : yes
Redirection Policy                  :
Per Host Replication                : no
Egress Rate Modify                  : no
Mcast Reporting Destination Name    :
Mcast Reporting Admin State         : Disabled
===============================================================================
```

Below is a command to view the (S,G)s that all subscribers are requesting. Since the system has only one subscriber, this example only shows one host.

```
*A:BNG-1> show router igmp hosts detail
===============================================================================
IGMP Host 10.0.0.24
===============================================================================
Oper Status      : Up         MacAddress     : 00:00:10:10:10:11
Oper version     : 3          Subscriber     : Subscriber-1
Num Groups       : 1          GrpItf         : group-int-1
Max Grps Till Now: 1          IGMP-Policy    : igmp-policy-1
PPPoE SessionId  : N/A        Next query time: 0d 00:01:52
FwdSvcId         : 1          Max Srcs Allow*: No Limit
Max Grps Allowed : No Limit   Max Grp Srcs A*: No Limit
-------------------------------------------------------------------------------
IGMP Group
-------------------------------------------------------------------------------
Group Address    : 239.255.1.1      Up Time    : 0d 00:02:46
Expires          : Not running    Mode       : Include
V1 Host Timer    : Not running    Type       : Dynamic
V2 Host Timer    : Not running    Compat Mode: IGMP Version 3
Redir.SvcId      : N/A            Redir.Intf : N/A
-------------------------------------------------------------
Source Address     Expires        Type         Fwd/Blk
-------------------------------------------------------------
```

```
192.168.4.2        0d 00:04:09    Dynamic      Fwd
-------------------------------------------------------------------------------
Hosts : 1
```

To check for an individual subscriber and its requested (S,G)s, the following command can be used.

```
*A:BNG-1> show service active-subscribers igmp subscriber "Subscriber-2" detail
===============================================================================
Active Subscribers Detail
===============================================================================
Subscriber                      IGMP-Policy
  HostAddr                      GrpItf                            NumGroups
    GrpAddr                     Type            Up-Time           Mode
      SrcAddr                   Type                              Blk/Fwd
-------------------------------------------------------------------------------
Subscriber-1                    igmp-policy-1
  10.0.0.24                     group-int-1                          1
    239.255.1.1                 Dynamic         0d 00:01:26       Include
      192.168.4.2               Dynamic                           Fwd
-------------------------------------------------------------------------------
Number of Subscribers : 1
===============================================================================
```

ESM IGMP PPPoE Walkthrough

IGMP message processing and delivery of multicast content for PPPoE subscribers is considered next. Figure 124 shows the message flow for multicast content delivery to PPPoE subscribers.

Figure 124 **PPPoE Multicast Flow**

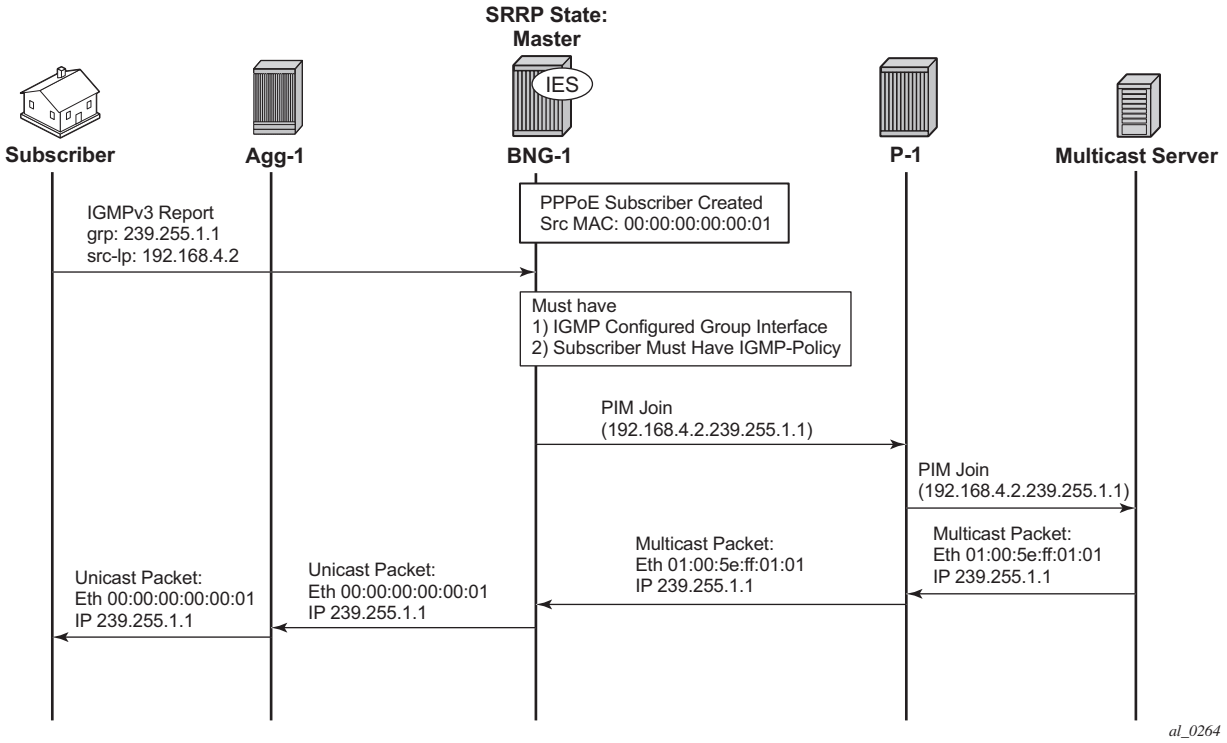

As stated earlier, the important configuration aspects are highlighted in the dotted box. The main difference between IPoE subscribers and PPPoE subscribers is the multicast data path. PPPoE subscribers receive multicast content via Ethernet unicast while IPoE subscribers receive multicast content via Ethernet multicast. PPPoE natively does not have a multicast mechanism and requires all data traffic to be unicasted. Even if the subscribers are on the same SAP, multicast content is replicated per subscriber. To achieve this, the IP header indicates a multicast address while the Ethernet header destination MAC address is changed to the subscriber's MAC address.

Step 1. Verify the group interface. It is very similar to the output for the IPoE group interface.

```
*A:BNG-1> show router igmp group-interface detail
===============================================================================
IGMP Group-Interfaces
===============================================================================
FwdSvc/Grp-Intf    : 1/group-int-1
Admin-Status       : Up                  Oper-Status        : Up
Import-Policy      : none                Subnet-Check       : Enabled
Router-Alert-Check : Enabled             Sub-Hosts-Only     : Enabled
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
-------------------------------------------------------------------------------
 SAP               : 1/1/5:4
 Admin/Oper version: 3/3                Num Groups         : 0
 Max Groups Allowed: No Limit           Max Groups Till Now: 0
 Max Sources Allow*: No Limit
 Max Grp Srcs Allo*: No Limit
-------------------------------------------------------------------------------
Group-Interfaces = 1, SAPs = 1
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

Next an IGMPv3 message is sent towards the BNG. The (S,G) is (192.168.4.2, 239.255.1.1). The PPPoE subscriber is assigned an IP address of 10.0.0.12.

The output below shows the key difference between a PPPoE subscriber and an IPoE subscriber. PPPoE multicast content is replicated per host and not per SAP. This output shows this clearly as the multicast group is associated with the host and not with the SAP.

```
*A:BNG-1> show router igmp group
===============================================================================
IGMP Interface Groups
===============================================================================
===============================================================================
IGMP Host Groups
===============================================================================
(192.168.4.2,239.255.1.1)
    Fwd List  : 10.0.0.12          Up Time : 0d 17:19:08
===============================================================================
IGMP SAP Groups
===============================================================================
-------------------------------------------------------------------------------
(*,G)/(S,G) Entries : 1
===============================================================================
```

The next command shows all of the subscribers and all the (S,G)s joined. In this case there is only one PPPoE subscriber.

```
*A:BNG-1> show router igmp hosts detail
===============================================================================
IGMP Host 10.0.0.12
===============================================================================
Oper Status      : Up          MacAddress     : 52:e0:50:bd:00:00
Oper version     : 3           Subscriber     : User-ppp-1
Num Groups       : 1           GrpItf         : group-int-1
Max Grps Till Now: 1           IGMP-Policy    : igmp-policy-1
PPPoE SessionId  : 1           Next query time: 0d 00:01:47
FwdSvcId         : 1           Max Srcs Allow*: No Limit
Max Grps Allowed : No Limit    Max Grp Srcs A*: No Limit

-------------------------------------------------------------------------------
```

```
IGMP Group
-------------------------------------------------------------------------------
Group Address    : 239.255.1.1    Up Time    : 0d 00:00:36
Expires          : Not running    Mode       : Include
V1 Host Timer    : Not running    Type       : Dynamic
V2 Host Timer    : Not running    Compat Mode: IGMP Version 3
Redir.SvcId      : N/A            Redir.Intf : N/A
-------------------------------------------------------------
Source Address     Expires        Type         Fwd/Blk
-------------------------------------------------------------
192.168.4.2        0d 00:04:03    Dynamic      Fwd
-------------------------------------------------------------------------------
Hosts : 1
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

To view each individual subscriber and their respective (S,G)s, use the command below.

```
*A:BNG-1> show service active-subscribers igmp subscriber "user02" detail
===============================================================================
Active Subscribers Detail
===============================================================================
Subscriber                      IGMP-Policy
  HostAddr                      GrpItf                          NumGroups
    GrpAddr                     Type            Up-Time         Mode
      SrcAddr                   Type                            Blk/Fwd
-------------------------------------------------------------------------------
User-ppp-1                      igmp-policy-1
  10.0.0.12                     group-int-1                      1
    239.255.1.1                 Dynamic         0d 00:02:07      Include
      192.168.4.2               Dynamic                          Fwd
-------------------------------------------------------------------------------
Number of Subscribers : 1
===============================================================================
```

ESM IGMP MCS

The BNGs are configured with SRRP for both IPoE and PPPoE subscribers. This provides stateful redundancy when the master BNG fails. The master BNG will be the only one processing and answering IGMP messages. The standby BNG does not perform any IGMP processing and receives updates through MCS for all subscribers in real time. In the event of a failure, the standby will become active and starts processing all IGMP messages. The standby will also immediately trigger PIM joins for all of the subscribers’s (S,G)s. This is all possible because the standby is always synchronized with the master BNG prior to the failover. Restoration of all multicast channels should happen quickly after the failover and depends on both the PIM configuration and the underlying routing infrastructure.

The key parameters for MCS for ESM multicast are: syncing of subscribers (ipoe, pppoe), SRRP and IGMP. Below is the redundancy configuration for BNG-1 and BNG-2.

```
*A:BNG-1>config>redundancy# info
----------------------------------------------
        multi-chassis
            peer 192.0.2.2 create
                sync
                    igmp
                    srrp
                    sub-mgmt ipoe pppoe
                    port 1/1/5 create
                        range 4-4 sync-tag "sub"
                        range 5-5 sync-tag "srrp"
                    exit
                    no shutdown
                exit
                no shutdown
            exit
        exit
```

The following command displays the number of entries being synced across the BNGs.

```
*A:BNG-1> show redundancy multi-chassis sync peer 192.0.2.2 detail
===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 192.0.2.2
Description             : (Not Specified)
Authentication          : Disabled
Source IP Address       : 192.0.2.1
Admin State             : Enabled
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     : IGMP SUBMGMT-IPOE SUBMGMT-PPPOE SRRP
Sync Admin State        : Up
Sync Oper State         : Up
DB Sync State           : inSync
Num Entries             : 15
Lcl Deleted Entries     : 0
Alarm Entries           : 0
Rem Num Entries         : 15
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
===============================================================================
MCS Application Stats
===============================================================================
Application             : igmp
Num Entries             : 1
Lcl Deleted Entries     : 0
Alarm Entries           : 0
```

```
-------------------------------------------------------------------------------
Rem Num Entries         : 1
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
-------------------------------------------------------------------------------
Application             : subMgmtIpoe
Num Entries             : 1
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 1
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
-------------------------------------------------------------------------------
Application             : srrp
Num Entries             : 14
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 14
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
-------------------------------------------------------------------------------
Application             : subMgmtPppoe
Num Entries             : 1
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 1
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
-------------------------------------------------------------------------------
===============================================================================
```

To check the details of the synchronized data across the BNGs, use the command below. It provides a detailed description of the IGMP information synced across MCS.

```
*A:BNG-1> tools dump redundancy multi-chassis sync-database application igmp detail

If no entries are present for an application, no detail will be displayed.

FLAGS LEGEND: ld - local delete; da - delete alarm; pd - pending global delete

Peer Ip 192.0.2.2


Application IGMP
Sap-id              Client Key
 SyncTag                          DLen  Flags    timeStamp
  deleteReason code and description
-------------------------------------------------------------------------------
1/1/5:4             Host=10.0.0.10, HostGroup=239.255.1.1
 sub                              20    -- -- -- 05/23/2013 13:05:31
  0x0

The following totals are for:
 peer ip ALL, port/lag ALL, sync-tag ALL, application IGMP
```

```
Valid Entries:                   1
Locally Deleted Entries:         0
Locally Deleted Alarmed Entries: 0
Pending Global Delete Entries:   0
```

ESM IGMP Debug

There are many debug features for ESM multicast. Debug allows real-time monitoring of all events happening on the system and can assist operators with troubleshooting. First enable debug on the system, then send an IGMP message to join a multicast group (S,G). Again the IGMP message used in this case is IGMPv3 with SSM. Below is the debug information for ESM IGMP at packet level.

```
debug
    router
        igmp
            packet mode egr-ingr-and-dropped
        exit
    exit

2977 2013/05/23 13:01:45.43 EST MINOR: DEBUG #2001 IGMP[9]
"IGMP[9]: RX-PKT
[012 03:58:58.090] IGMP host 10.0.0.10 V3 PDU: 10.0.0.10 -> 224.0.0.22 pduLen
20
    Type: V3 REPORT maxrespCode 0x0 checkSum 0xddf7
    Num Group Records: 1
        Group Record 0
        Type: ALW_NEW_SRCS, AuxDataLen 0, Num Sources 1
        Mcast Addr: 239.255.1.1
        Source Address List
            192.168.4.2

"
```

Below is the debug information for ESM IGMP at host level and the associated IGMP events.

```
debug
    router
        igmp
            host "10.0.0.10"
        exit
    exit

2978 2013/05/23 13:01:45.43 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpIfGroupAdd
Adding 239.255.1.1 to IGMP host 10.0.0.10 database"

2979 2013/05/23 13:01:45.43 EST MINOR: DEBUG #2001 ies1 IGMP[ies1  inst 9
]
```

```
"IGMP[ies1 inst 9]: igmpProcessGroupRec
Process group rec ALW_NEW_SRCS received on host 10.0.0.10 for group 239.255.1.1 i
n mode INCLUDE. Num srcs 1"

2980 2013/05/23 13:01:45.43 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpIfSrcAdd
Adding i/f source entry for host 10.0.0.10 (192.168.4.2,239.255.1.1) to IGMP fwdList
 Database, redir if N/A"
```

Below is the debug information for ESM IGMP if MCS synchronization is enabled.

```
debug
    router
        igmp
            mcs "group-int-1"
        exit
    exit


2981 2013/05/23 13:01:45.44 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpMcsAddIfGroup
Building MCS entry for host 10.0.0.10, group 239.255.1.1"
```

The same debug commands can be used for viewing subscribers IGMP leave messages. Below is the debug information for ESM IGMP at packet level.

```
debug
    router
        igmp
            packet mode egr-ingr-and-dropped
        exit
    exit

2982 2013/05/23 13:02:23.75 EST MINOR: DEBUG #2001 ies1 IGMP[9]
"IGMP[9]: RX-PKT
[012 03:59:36.410] IGMP host 10.0.0.10 V3 PDU: 10.0.0.10 -> 224.0.0.22 pduLen
20
    Type: V3 REPORT maxrespCode 0x0 checkSum 0xdcf7
    Num Group Records: 1
        Group Record 0
        Type: BLK_OLD_SRCS, AuxDataLen 0, Num Sources 1
        Mcast Addr: 239.255.1.1
        Source Address List
            192.168.4.2
```

Below is the debug information for ESM IGMP at host level and the associated IGMP events.

```
debug
```

```
    router
        igmp
            host "10.0.0.10"
        exit
    exit

2983 2013/05/23 13:02:23.75 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpProcessGroupRec
Process group rec BLK_OLD_SRCS received on host 10.0.0.10 for group 239.255.1.1 i
n mode INCLUDE. Num srcs 1"

2984 2013/05/23 13:02:23.75 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpProcessIfSrcTimerExp
Source Timer expired for IGMP host 10.0.0.10 (192.168.4.2,239.255.1.1)"

2985 2013/05/23 13:02:23.75 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpIfSrcDel
Deleting i/f source entry for host 10.0.0.10 (192.168.4.2,239.255.1.1) from IGMP Dat
abase. DeleteFromAvl: 1 !Redir 0"

2986 2013/05/23 13:02:23.75 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpIfGroupDel
Deleting 239.255.1.1 from IGMP host 10.0.0.10 database"
```

The debug information when MCS removes the entry on the standby BNG is shown below.

```
debug
    router
        igmp
            mcs "group-int-1"
        exit
    exit
2987 2013/05/23 13:02:23.75 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpMcsDelIfGroup
Deleting MCS entry for host 10.0.0.10, group 239.255.1.1, Glb"

2988 2013/05/23 13:02:23.75 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpMcsDelIfGroup
Deleting MCS entry for host 10.0.0.10, group 239.255.1.1, Glb"

2989 2013/05/23 13:02:23.75 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpMcsDelIfGroup
Deleting MCS entry for host 10.0.0.10, group 239.255.1.1, Glb"
```

IGMP Control Plane Filters

IGMP control plane filtering can be applied at the router level and/or subscriber level (IGMP-policy). The filter list contains multicast groups (S,G) and is provisioned at the router level in the policy-options context. The filter can be applied either as a black-list or a white-list.

Step 1. Provision a prefix list for the multicast group (G). The configuration below is an example showing the various options possible for the prefix list. The only one used in this configuration is the prefix 239.255.1.1/32.

```
*A:BNG-1> config>router>policy-options#
            prefix-list "igmp-prefix-list-1"
                prefix 239.255.1.1/32 exact
                prefix 239.255.2.0/24 longer
                prefix 239.255.3.0/24 prefix-length-range 24-25
                prefix 239.255.4.0/24 through 25
            exit
```

Step 2. Provision a router policy.

a. Step 2a: Option 1 white-list, used below.

Provision a router policy. In the example below a white-list is configured, allowing only the prefix list specified and reject everything else. Source-address configuration is also possible for IGMP v3 (S,G).

```
*A:BNG-1> config>router>policy-options#
            policy-statement "igmp-white-list-1"
                entry 10
                    from
                        group-address "igmp-prefix-list-1"
                        source-address 192.168.4.2
                    exit
                    action accept
                    exit
                exit
                default-action reject
            exit
```

b. Step 2b: Option 2 black-list, not used below.

Provision a router policy. Here, a black-list is configured, denying the prefix list and accepting everything else. Again, source-address configuration is also possible for IGMP v3 (S,G).

```
*A:BNG-1> config>router>policy-options#
            policy-statement "igmp-black-list-1"
                entry 10
                    from
                        group-address "igmp-prefix-list-1"
                        source-address 192.168.4.2
```

```
                exit
                action reject
                exit
            exit
            default-action accept
        exit
```

Step 3. Apply a Hierarchical filter.

a. Step 3a: Hierarchical filter, group-interface.

The filter can be applied in two places. First, at router/group-interface level, this will apply to all subscribers connected to the group interface.

```
*A:BNG-1> config>router
            igmp
                group-interface "group-int-1"
                    import "igmp-white-list-1"
                    no shutdown
                exit
                no shutdown
            exit
```

b. Step 3b: Hierarchical filter, subscriber level.

The group-interface filter takes precedence over the subscriber level filter. After the group-interface applies its filter against the incoming IGMP messages, the individual subscriber defined IGMP filters will be applied to the remaining IGMP messages.

```
*A:BNG-1> config>subscr-mgmt>igmp-policy# info
----------------------------------------------
            import "igmp-white-list-1"
```

Use the **debug** command to verify that the policy is performing correctly for the host. Note that group 239.255.1.2 is not in the white-list and so is dropped.

```
debug
    router
        igmp
            group-interface "group-int-1"
            host "10.0.0.10"
            packet mode egr-ingr-and-dropped
            mcs
        exit
    exit


3310 2013/05/23 14:51:55.69 EST MINOR: DEBUG #2001 ies1 IGMP[9]
"IGMP[9]: RX-PKT
[012 05:49:08.350] IGMP host 10.0.0.10 V3 PDU: 10.0.0.10 -> 224.0.0.22 pduLen
20
    Type: V3 REPORT maxrespCode 0x0 checkSum 0xe1f6
    Num Group Records: 1
        Group Record 0
```

```
        Type: MODE_IS_INCL, AuxDataLen 0, Num Sources 1
        Mcast Addr: 239.255.1.2
        Source Address List
            192.168.4.2

"

3311 2013/05/23 14:51:55.69 EST MINOR: DEBUG #2001 ies1 IGMP[ies1 inst 9
]
"IGMP[ies1 inst 9]: igmpParseV3Report
IGMP V3 policy DROP on host 10.0.0.10, from host 10.0.0.10, grpAddr 239.255.1.2,
 srcAddr 192.168.4.2"
```

IGMP Data Plane Filters

IGMP data plane filter utilize the ip-filter defined in the sla-profile. Again the filter can be used as a black-list or a white-list.

Step 1. Configure an ip-filter. Below is an example of a black-list filter.

```
*A:BNG-1> config>filter>ip-filter$ info
----------------------------------------------
            default-action forward
            entry 1 create
                match
                    dst-ip 239.255.1.1/32
                exit
                action drop
            exit
```

Step 2. Apply the configured ip filter into an sla-profile. Since multicast content is sent towards the subscriber, it is applied to the sla-profile egress.

```
*A:BNG-1> config>subscr-mgmt>sla-prof# info
----------------------------------------------
            egress
                ip-filter 1
                exit
            exit
----------------------------------------------
```

To view the statistics of the filter applied to the subscribers use the following command.

```
*A:BNG-1> show filter ip 1 counters
===============================================================================
IP Filter
===============================================================================
Filter Id    : 1                                Applied        : Yes
Scope        : Template                         Def. Action    : Forward
```

```
Radius Ins Pt: n/a
CrCtl. Ins Pt: n/a
RadSh. Ins Pt: n/a
Entries      : 1
Description  : (Not Specified)
-------------------------------------------------------------------------------
Filter Match Criteria : IP
-------------------------------------------------------------------------------
Entry        : 1
Ing. Matches : 0 pkts
Egr. Matches : 0 pkts
===============================================================================
```

Conclusion

Multicast is an essential part of Triple Play Services. The 7750 SR TPSDA solution offers much more than a baseline multicast delivery, it includes individual subscriber awareness and a full state redundancy option. Subscriber awareness allows for the fine tuning of each subscriber's multicast experience and also for troubleshooting on a per subscriber basis. Full state redundancy reduces failover time and ensures high availability of the services offered. This example provides a complete configuration walkthrough of both the IPoE and PPPoE SRRP models.

For operators wanting to further control and restrict individual subscriber's multicast content, ESM has a comprehensive set of both control path filtering and data path filtering.

ESM SLAAC Prefix Assignment via Local Address Server

This chapter provides information about ESM SLAAC prefix assignment via local address server.

Topics in this chapter include:

Applicability

This chapter is applicable to the 7750 SR-12 with Input/Output Module IOM3-XP and Integrated Media Module (IMM), and requires chassis mode C as a minimum. This feature is also supported on the 7450 ESS chassis in mixed-mode and the 7750 SR-c4/7/12 platform.

The configuration was tested on release 13.0.R1 and supports both Internet Protocol over Ethernet (IPoE) and Point-to-Point Protocol over Ethernet (PPPoE) subscribers.

Overview

Triple Play Service Delivery Architecture (TPSDA) supports IPv6 address/prefix assignment through Dynamic Host Configuration Protocol (DHCP), Point-to-Point Protocol (PPP), and Stateless Address Auto-Configuration (SLAAC). This chapter provides a configuration example of SLAAC prefix assignment via the local address server.

The network topology shown in Figure 125 shows a TPSDA setup. The setup consists of a 7750 SR serving as a Broadband Network Gateway (BNG). The 7450 is used as a Layer 2 switch aggregating all subscriber traffic.

Figure 125 **TPSDA Network Topology**

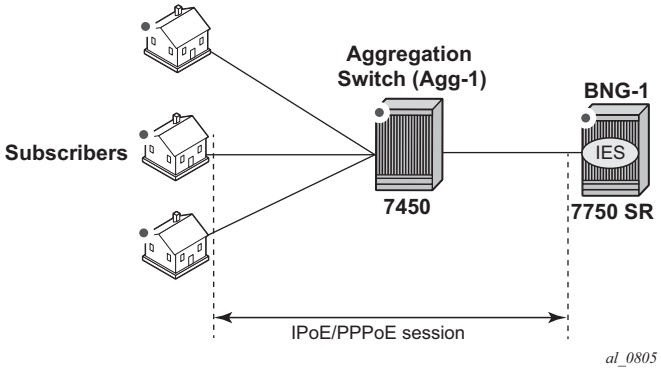

There are two methods available for subscriber SLAAC prefix assignment. The first method, not covered in this chapter, is to pre-define a static SLAAC prefix for each subscriber on the BNG, in a Local User Database (LUDB) or via a RADIUS AAA server. With such a configuration, the database would contain hundreds of thousands of /64 SLAAC prefixes, each with their associated host. Every time a subscriber moved to a new location (a new subnet), the allocation of a new prefix within the new subnet would be required, followed by a manual update of the database.

The second method, covered in this chapter, is to simplify SLAAC prefix assignment. A local address server is configured to dynamically assign SLAAC prefixes to hosts. Only a SLAAC pool name is obtained from RADIUS or LUDB after a successful subscriber authentication. This pool name is then used to assign a SLAAC prefix to the subscriber, out of the named address pool.

Using a local pool for SLAAC prefix assignment provides the following advantages:

- Reduces the configuration required on the RADIUS server, the local user database (LUDB), and the BNG to a few lines;
- Removes the complexity of managing actual prefixes in a database;
- Reduces configuration errors; for example, accidentally assigning the same prefix to two different subscribers.

The local address server already has tools, logs, and monitoring features for prefix management, such as prefix depletion and subnet migration. Service providers can rely on the local address server to assist in SLAAC prefix assignment.

Configuration

This guide assumes a basic knowledge of ESM.

Different Types of SLAAC Hosts

SLAAC is supported for both PPP and IP over Ethernet hosts. The local address server can be enabled for either or both host types on a group interface level.

PPP SLAAC Hosts

PPP IPv4 hosts rely on IPCP to retrieve an IPv4 address. However, IPv6CP does not assign IPv6 addresses or prefixes to the host. PPP hosts rely on router solicitations (RSs) or DHCPv6 to obtain IPv6 addresses/prefixes.

PPP SLAAC hosts creation requires three configuration steps.

Step 1. Following is a baseline PPP subscriber management configuration on the BNG.

```
*A:BNG-1>config>service>ies# info
----------------------------------------------
            description "BNG-1"
            subscriber-interface "sub-int-1" create
                ipv6
                    subscriber-prefixes
                        prefix 2001:db8::/32 wan-host
                    exit
                exit
                group-interface "group-int-1" create
                    ipv6
                        router-advertisements
                            prefix-options
                                autonomous
                            exit
                            no shutdown
                        exit
                    exit
                    sap 1/1/5:4 create
                        sub-sla-mgmt
                            def-sub-id use-sap-id
                            def-sub-profile "sub-profile-1"
                            def-sla-profile "sla-profile-1"
                            sub-ident-policy "sub-ident-policy-1"
                            multi-sub-sap 10
                            no shutdown
```

```
                    exit
                exit
                pppoe
                    no shutdown
                exit
            exit
        exit
```

Step 2. A DHCPv6 server is used as the local address server to assign SLAAC prefixes. It is possible to reuse the same pool for both DHCPv6 and SLAAC address/prefix assignment. For SLAAC hosts, the keyword wan-host is required.

```
*A:P-1>config>router>info
#------------------------------------------------
echo "Local DHCP Server Configuration"
#------------------------------------------------
          dhcp6
                local-dhcp-server "dhcp6-server-1" create
                    use-pool-from-client
                    pool "pool-v6-1" create
                        prefix 2001:db8::/32 wan-host create
                        exit
                    exit
```

Step 3. On the PPP group interface, configure the local address server. Specify that the local address server is to be used for client application ppp-slaac. The server name must match the name configured for the DHCPv6 server (step 2). The DHCPv6 server is reused as the local address server.

```
*A:BNG-1>config>service>ies>sub-if>grp-if# info
----------------------------------------------
                    local-address-assignment
                        ipv6
                            client-application ppp-slaac
                            server "dhcp6-server-1"
                        exit
                        no shutdown
                    exit
                exit
            exit
```

There are two options for supplying a SLAAC pool name for a PPP host: RADIUS and LUDB.

Option 1: During PPP authentication, RADIUS can return the SLAAC pool name attribute along with other subscriber attributes. Note: Remove the user-db configuration from pppoe when using RADIUS authentication. Add an authentication policy to the group interface to enable RADIUS authentication.

```
*A:BNG-1>config>service>ies>sub-if>grp-if# info
----------------------------------------------
```

```
                    authentication-policy "auth-policy-1"
                exit
            exit
```

Then, add the attribute Alc-SLAAC-IPv6-Pool to the RADIUS database. The following is an example from a freeradius clients file.

```
user_ppp_01  Auth-Type := CHAP,    Cleartext-Password := password
       Alc-SLA-Prof-Str = "sla-profile-1",
       Alc-Subsc-ID-Str = "home-ppp-1",
       Alc-Subsc-Prof-Str = "sub-profile-1",
       Alc-SLAAC-IPv6-Pool = pool-v6-1,
       Alc-PPP-Force-IPv6CP = 1
```

Option 2: During PPP authentication, an LUDB can return the SLAAC pool name attribute along with other subscriber attributes. Note: Remove the authentication policy from the group interface when using LUDB.

First, create an LUDB and add a user to the LUDB. This LUDB is configured with a default host for all PPPoE hosts and returns a default SLAAC pool name.

```
*A:BNG-1>config>subscr-mgmt# info
----------------------------------------------
            local-user-db "pppoe-ludb-lookup" create
            ppp
                match-list username
                host "default" create
                    ipv6-slaac-prefix-pool "pool-v6-1"
                    no shutdown
                exit
            exit
            no shutdown
        exit
```

Then, reference this LUDB from the group interface.

```
*A:BNG-1>config>service>ies>sub-if>grp-if# info
----------------------------------------------
                  pppoe
                       user-db “pppoe-ludb-lookup”
                       no shutdown
                  exit
            exit
```

With the preceding configuration, this group interface supports SLAAC prefix assignment through the local address pool. The following is the result of a PPP host being assigned a SLAAC prefix by the local address server. In the **show pppoe session**, the IPv6 prefix is from the local address pool and the pool name is from authentication (RADIUS or LUDB).

```
*A:BNG-1> show  service active-subscribers hierarchy
===============================================================================
Active Subscriber hierarchy
```

```
===============================================================================
-- 00:00:64:19:01:03|1/1/5:4|1 (sub-profile-1)
   |
   |-- sap:1/1/5:4 - sla:sla-profile-1
   |   |
   |   |-- 2001:db8::/64
   |   |   00:00:64:19:01:03 - 1 (PPP-SLAAC)
   |   |


*A:BNG-1> show  service id 1 pppoe session detail
===============================================================================
PPPoE sessions for svc-id 1
===============================================================================
Sap Id               Mac Address       Sid    Up Time         Type
    IP/L2TP-Id/Interface-Id                                      MC-Stdby
-------------------------------------------------------------------------------
1/1/5:4              00:00:64:19:01:03 1      0d 00:02:06     local
    02:00:64:FF:FE:19:01:03

LCP State            : Opened
IPCP State           : Closed
IPv6CP State         : Opened
PPP MTU              : 1492
PPP Auth-Protocol    : None
PPP User-Name        : (Not Specified)

Subscriber-interface : sub-int-1
Group-interface      : grp-int-1

IP Origin            : none
DNS Origin           : none
NBNS Origin          : none

Subscriber           : "00:00:64:19:01:03|1/1/5:4|1"
Sub-Profile-String   : ""
SLA-Profile-String   : ""
ANCP-String          : ""
Int-Dest-Id          : ""
App-Profile-String   : ""
Category-Map-Name    : ""
Acct-Session-Id      : "D84FFF0000005A555513BC"
Sap-Session-Index    : 1

IP Address           : N/A
Primary DNS          : N/A
Secondary DNS        : N/A
Primary NBNS         : N/A
Secondary NBNS       : N/A
Address-Pool         : N/A

IPv6 Prefix          : 2001:db8::/64
IPv6 Prefix Origin   : local-pool
IPv6 Prefix Pool     : "pool-v6-1"
IPv6 Del.Pfx.        : N/A
IPv6 Del.Pfx. Origin : none
IPv6 Del.Pfx. Pool   : ""
IPv6 Address         : N/A
IPv6 Address Origin  : none
```

```
IPv6 Address Pool    : ""
Primary IPv6 DNS     : N/A
Secondary IPv6 DNS   : N/A

Circuit-Id           : circuit1
Remote-Id            :

Radius Session-TO    : N/A
Radius Class         :
Radius User-Name     :
Logical-Line-Id      :
Service-Name         : AGILENT
-------------------------------------------------------------------------------
Number of sessions   : 1
===============================================================================
```

ICMP6 debugging can be used to show the SLAAC address assignment process.

```
debug
    router
        ip
            icmp6
        exit
    exit

23673 2015/05/14 21:38:38.04 UTC MINOR: DEBUG #2001 TIP
"TIP: ICMP6_PKT
ICMP6 egressing on grp-int-01 (ies1):
   fe80::101 -> ff02::1
   Type: Router Advertisement (134)
   Code: No Code (0)
      Hop Limit     : 64
      Flags         :
      Retrans Time  : 0
      Def Life Time : 4500
      Reachable Time: 0
      Option  : Prefix         : 2001:db8::/64
                Flags          : On Link Autoconfig
                Valid Life Time: 86400
                Pref  Life Time: 3600

"
```

IPoE SLAAC Hosts

IPoE offers two methods to create an SLAAC host:

1. Triggered by a successful IPv4 host creation
2. Triggered by an RS request

SLAAC Host Creation via IPv4 Host

A successful IPv4 host creation can subsequently trigger the creation of a SLAAC host; this is known as IPoE-linking. The SLAAC prefix for the host must be provided through either RADIUS or LUDB during the IPv4 host authentication.

IPoE SLAAC host creation through IPoE linking requires four steps.

Step 1. Following is a baseline IPoE subscriber management configuration on the BNG.

```
*A:BNG-1>config>service>ies# info
----------------------------------------------
            description "BNG-1"
            subscriber-interface "sub-int-1" create
                address 10.255.255.253/8
                ipv6
                    subscriber-prefixes
                        prefix 2001:db8::/32 wan-host
                    exit
                exit
                group-interface "group-int-1" create
                    dhcp
                        server 192.168.0.1
                        lease-populate 10
                        client-applications dhcp
                        gi-address 10.255.255.253
                        no shutdown
                    exit
                    ipv6
                        router-advertisements
                            prefix-options
                                autonomous
                            exit
                            no shutdown
                        exit
                    exit
                    sap 1/1/5:4 create
                        sub-sla-mgmt
                            def-sub-id use-sap-id
                            def-sub-profile "sub-profile-1"
                            def-sla-profile "sla-profile-1"
                            sub-ident-policy "sub-ident-policy-1"
                            multi-sub-sap 10
                            no shutdown
                        exit
                    exit
                exit
            exit
```

Step 2. Enable IPoE-linking to allow SLAAC host creation after a successful IPv4 host creation. Several options should be enabled for the SLAAC host to function. Gratuitous router advertisement will send unsolicited router advertisements with a SLAAC prefix for the host to use. The BNG uses the gratuitous router advertisement to let the subscriber know the assigned prefix to auto-configure. In this case, where prefixes are dynamically assigned, the subscriber will not know the prefix ahead of time, so the gratuitous router advertisement must be enabled. Shared-circuit-id will allow the SLAAC host to use the same circuit ID as the IPv4 host.

```
*A:BNG-1>config>service>ies>sub-if>grp-if# info
----------------------------------------------
                    ipoe-linking
                        shared-circuit-id
                        gratuitous-rtr-adv
                        no shutdown
                    exit
            exit
```

Step 3. As with PPP hosts, the DHCPv6 server is reused as the local address server for SLAAC prefix assignment. It is possible to reuse the same pool for both DHCPv6 and SLAAC subscribers. For SLAAC hosts, the keyword wan-host is required. In this case, an IPv4 host must be created first to trigger the creation of the IPv6 SLAAC host. The following example uses the local DHCPv4 server for IPv4 address assignment, but it is possible to use other methods for IPv4 address assignment, such as through LUDB and RADIUS proxy.

```
*A:P-1>config>router>info
#------------------------------------------------
echo "Local DHCP Server Configuration"
#------------------------------------------------
          dhcp
            local-dhcp-server "dhcp-server-1" create
                use-gi-address scope pool
                pool "pool-v4-1" create
                    subnet 10.0.0.0/8 create
                        options
                            subnet-mask 255.0.0.0
                            default-router 10.255.255.253
                        exit
                        address-range 10.0.0.10 10.0.0.254
                    exit
                exit
                no shutdown
          exit
          dhcp6
                local-dhcp-server "dhcp6-server-1" create
                    use-pool-from-client
                    pool "pool-v6-1" create
                        2001:db8::/32 wan-host create
                        exit
                    exit
                exit
                no shutdown
```

```
            exit
```

Step 4. On the group interface, configure the local address server. Specify that the local address-server is to be used for client application ipoe-slaac. The server name must match the name configured for the DHCPv6 server (Step 3). The local address server reuses the local DHCPv6 server.

```
*A:BNG-1>config>service>ies>sub-if>grp-if# info
----------------------------------------------
                    local-address-assignment
                        ipv6
                            client-application ipoe-slaac
                            server "dhcp6-server-1"
                        exit
                        no shutdown
                    exit
```

There are two options for supplying a SLAAC pool name for the DHCPv4 host: RADIUS and LUDB.

Option 1: During authentication, RADIUS can return the SLAAC pool name attribute along with other subscriber attributes. Note: Remove the user-db configuration from the DHCP and IPOE-session context when using RADIUS authentication. First, add an authentication policy to the group interface to allow RADIUS authentication.

```
*A:BNG-1>config>service>ies>sub-if>grp-if# info
----------------------------------------------
                    authentication-policy "auth-policy-1"
                exit
            exit
```

Then, add the attribute Alc-SLAAC-IPv6-Pool to the subscriber host RADIUS user database. The following is an example using the client file on freeradius.

```
00:00:10:10:12:13  Cleartext-Password := password
       Alc-SLA-Prof-Str = "sla-profile-1",
       Alc-Subsc-ID-Str = "home-ipoe-1",
       Alc-Subsc-Prof-Str = "sub-profile-1",
       Alc-SLAAC-IPv6-Pool = pool-v6-1
```

Option 2: During authentication, LUDB can return the SLAAC pool name attribute along with other subscriber attributes. Note: Remove the authentication policy from the group interface when using LUDB.

First, create an LUDB and add a user in the LUDB. This LUDB is configured with a default host for all DHCPv4 hosts and returns a default SLAAC pool name.

```
*A:BNG-1>config>subscr-mgmt# info
----------------------------------------------
        local-user-db "ipoe-ludb-lookup" create
            ipoe
                match-list sap-id
```

```
                host "default" create
                    ipv6-slaac-prefix-pool "pool-v6-1"
                    no shutdown
                exit
            exit
            no shutdown
        exit
```

Then, reference this LUDB from the group interface. The LUDB can be referenced in two places.

Nokia recommends that IPoE subscribers use an IPoE session. In this case, the LUDB is referenced from the group interface ipoe-session context.

```
*A:BNG-1>config>service>ies>sub-if>grp-if# info
----------------------------------------------
                    ipoe-session
                         user-db "ipoe-ludb-lookup"
                         no shutdown
                    exit
```

For operators that do not use IPoE sessions (not recommended), the LUDB is referenced from the group interface dhcp context.

```
*A:BNG-1>config>service>ies>sub-if>grp-if# info
----------------------------------------------
                    dhcp
                         user-db "ipoe-ludb-lookup"
                         no shutdown
                    exit
```

With the preceding configuration, the local address server on the group interface is ready to assign SLAAC prefixes. Start a DHCPv4 session to the group interface SAP. In the show ipoe session, the IPv6 prefix origin is from the local address pool and the pool name is from authentication (RADIUS or LUDB).

```
*A:BNG-1> show service active-subscribers hierarchy
===============================================================================
Active Subscriber hierarchy
===============================================================================
-- 00:00:64:19:01:03|1/1/5:4|1 (sub-profile-1)
   |
   |-- sap:1/1/5:4 - sla:sla-profile-1
   |   |
   |   |-- 10.0.0.2
   |   |   00:00:64:19:01:03 - N/A (DHCP)
   |   |
   |   |-- 2001:db8::/64
   |   |   00:00:64:19:01:03 - N/A (IPoE-SLAAC)
   |   |


*A:BNG-1> show  service id 1 ipoe session detail
===============================================================================
```

```
IPoE sessions for service 1
===============================================================================

SAP                     : 1/1/5:4
Mac Address             : 00:00:64:19:01:03
Circuit-Id              : circuit12
Remote-Id               : remote0
Session Key             : sap-mac

MC-Standby              : No

Subscriber-interface    : sub-int-01
Group-interface         : grp-int-01

Up Time                 : 0d 00:07:02
Session Time Left       : N/A
Last Auth Time          : 05/15/2015 20:16:26
Min Auth Intvl (left)   : infinite (N/A)
Persistence Key         : N/A

Subscriber              : "00:00:64:19:01:03 |1/1/20:791"
Sub-Profile-String      : ""
SLA-Profile-String      : ""
ANCP-String             : ""
Int-Dest-Id             : ""
App-Profile-String      : ""
Category-Map-Name       : ""
Acct-Session-Id         : "D84FFF000000815556541A"
Sap-Session-Index       : 1

IP Address              : 10.0.0.2/8
IP Origin               : DHCP
Primary DNS             : N/A
Secondary DNS           : N/A
Primary NBNS            : N/A
Secondary NBNS          : N/A
Address-Pool            : N/A

IPv6 Prefix             : 2001:db8::/64
IPv6 Prefix Origin      : LclPool
IPv6 Prefix Pool        : "pool-v6-1"
IPv6 Del.Pfx.           : N/A
IPv6 Del.Pfx. Origin    : None
IPv6 Del.Pfx. Pool      : ""
IPv6 Address            : N/A
IPv6 Address Origin     : None
IPv6 Address Pool       : ""
Primary IPv6 DNS        : N/A
Secondary IPv6 DNS      : N/A

Radius Session-TO       : N/A
Radius Class            :
Radius User-Name        :
-------------------------------------------------------------------------------
Number of sessions : 1
===============================================================================
```

ICMP6 debugging can be used to show the SLAAC address assignment process.

```
debug
    router
        ip
            icmp6
        exit
    exit

23673 2015/05/14 21:38:38.04 UTC MINOR: DEBUG #2001 TIP
"TIP: ICMP6_PKT
ICMP6 egressing on grp-int-01 (ies1):
   fe80::101 -> ff02::1
   Type: Router Advertisement (134)
   Code: No Code (0)
      Hop Limit     : 64
      Flags         :
      Retrans Time  : 0
      Def Life Time : 4500
      Reachable Time: 0
      Option  : Prefix         : 2001:db8::/64
                Flags          : On Link Autoconfig
                Valid Life Time: 86400
                Pref  Life Time: 3600

"
```

SLAAC Host Creation via RS Trigger

An IPv6 SLAAC host can be created through a host router originated solicit message, which removes the dependency of a SLAAC host on successful DHCPv4 host creation.

SLAAC hosts creation via RS trigger requires four configuration steps.

Step 1. The following is a baseline IPoE subscriber management configuration on the BNG.

```
*A:BNG-1>config>service>ies# info
----------------------------------------------
            description "BNG-1"
            subscriber-interface "sub-int-1" create
                ipv6
                    subscriber-prefixes
                        prefix 2001:db8::/32 pd wan-host
                    exit
                exit
                group-interface "group-int-1" create
                    ipv6
                        router-advertisements
                            prefix-options
                                autonomous
                            exit
                            no shutdown
                        exit
                    exit
```

```
                    sap 1/1/5:4 create
                        sub-sla-mgmt
                            def-sub-id use-sap-id
                            def-sub-profile "sub-profile-1"
                            def-sla-profile "sla-profile-1"
                            sub-ident-policy "sub-ident-policy-1"
                            multi-sub-sap 10
                            no shutdown
                        exit
                    exit
                exit
            exit
```

Step 2. Enable the group interface to process router solicit messages. There are a few options available for router solicit triggered hosts. The inactivity timer will remove the host if the global unique address of the host is not learned through Neighbor Solicitation (NS), Router Solicitation (RS), or Duplicate Address Detection (DAD) messages within the time specified. The min-auth-interval is the interval that a subscriber must wait before the next router-solicit messages is used for re-authentication. Re-authentication can occur if the first RS was lost, or the BNG/RADIUS system was queued up with requests.

```
*A:BNG-1>config>service>ies>sub-if>grp-if>ipv6# info
----------------------------------------------
                    router-solicit
                            inactivity-timer min 5
                            min-auth-interval min 5
                            no shutdown
                    exit
```

Step 3. A DHCPv6 server is used as the local address server for SLAAC prefix assignment. It is possible to reuse the same pool for both DHCPv6 and SLAAC subscribers. For SLAAC hosts, the keyword **wan-host** is required.

```
*A:P-1>config>router>info
#------------------------------------------------
echo "Local DHCP Server Configuration"
#------------------------------------------------
          dhcp6
                local-dhcp-server "dhcp6-server-1" create
                    use-pool-from-client
                    pool "pool-v6-1" create
                        2001:db8::/32 wan-host create
                        exit
                    exit
                exit
                no shutdown
          exit
```

Step 4. On the group interface, configure the local address server. Specify that the local address server is to be used for client application ipoe-slaac. The server name must match the name configured for the DHCPv6 server. The local address server reuses the local DHCPv6 server.

```
*A:BNG-1>config>service>ies>sub-if>grp-if# info
----------------------------------------------
                    local-address-assignment
                        ipv6
                            client-application ipoe-slaac
                            server "dhcp6-server-1"
                        exit
                        no shutdown
                    exit
```

There are two options for supplying the SLAAC pool name for the DHCPv4 host: RADIUS and LUDB.

Option 1: During authentication, RADIUS can return the SLAAC pool name attribute along with other subscriber attributes. Note: Remove the user-db configuration from the router-solicit and ipoe-session when using RADIUS authentication. First, add an authentication policy to the group interface to allow RADIUS authentication.

```
*A:BNG-1>config>service>ies>sub-if>grp-if# info
----------------------------------------------
                    authentication-policy "auth-policy-1"
                exit
            exit
```

Then, add the attribute Alc-SLAAC-IPv6-Pool to the subscriber host RADIUS user database.

```
00:00:10:10:12:13  Cleartext-Password := password
       Alc-SLA-Prof-Str = "sla-profile-1",
       Alc-Subsc-ID-Str = "home-ipoe-1",
       Alc-Subsc-Prof-Str = "sub-profile-1",
       Alc-SLAAC-IPv6-Pool = pool-v6-1,
```

Option 2: During authentication, LUDB can return the SLAAC pool name attribute along with other subscriber attributes. Note: Remove the authentication policy from the group interface when using LUDB.

First, create an LUDB and add a user in the LUDB. The LUDB configures a default host for all SLAAC hosts and returns a default SLAAC pool name.

```
*A:BNG-1>config>subscr-mgmt# info
------------------------------------------------
        local-user-db "ipoe-ludb-lookup" create
            ipoe
                match-list username
                host "default" create
                    ipv6-slaac-prefix-pool "pool-v6-1"
                    no shutdown
                exit
            exit
            no shutdown
        exit
```

Then, reference an LUDB from the group interface. The LUDB can be referenced in two places.

Nokia recommends that IPoE subscribers use an IPoE session. In this case, the LUDB is referenced from the group interface ipoe-session context.

```
*A:BNG-1>config>service>ies>sub-if>grp-if# info
----------------------------------------------
                    ipoe-session
                         user-db "ipoe-ludb-lookup"
                         no shutdown
                    exit
```

For operators that do not enable IPoE sessions on the BNG (not recommended), the LUDB can be referenced from the group interface in the router-solicit context.

```
*A:BNG-1>config>service>ies>sub-if>grp-if>ipv6# info
----------------------------------------------
                    router-solicit
                         user-db "ipoe-ludb-lookup"
                         no shutdown
                    exit
```

With the preceding configuration, the group interface is ready to assign SLAAC prefixes from the local address pool. Let the host trigger a router solicit packet. In the show ipoe session, the IPv6 prefix is from the local address pool and the pool name is from authentication (RADIUS or LUDB).

```
*A:BNG-1> show service active-subscribers hierarchy
===============================================================================
Active Subscriber hierarchy
===============================================================================
-- 00:00:64:19:01:03|1/1/5:4|1 (sub-profile-1)
   |
   |-- sap:1/1/5:4 - sla:sla-profile-1
   |   |
   |   |-- 2001:db8::/64
   |   |   00:00:64:19:01:03 - N/A (IPoE-SLAAC)
   |   |
```

```
*A:BNG-1> show  service id 1 ipoe session detail
===============================================================================
IPoE sessions for service 1
===============================================================================
SAP                     : 1/1/5:4
Mac Address             : 00:00:64:19:01:03
Circuit-Id              : circuit12
Remote-Id               : remote0
Session Key             : sap-mac

MC-Standby              : No

Subscriber-interface    : sub-int-01
Group-interface         : grp-int-01
```

```
Up Time                 : 0d 00:07:02
Session Time Left       : N/A
Last Auth Time          : 05/15/2015 20:16:26
Min Auth Intvl (left)   : infinite (N/A)
Persistence Key         : N/A

Subscriber              : "00:00:64:19:01:03 |1/1/20:791"
Sub-Profile-String      : ""
SLA-Profile-String      : ""
ANCP-String             : ""
Int-Dest-Id             : ""
App-Profile-String      : ""
Category-Map-Name       : ""
Acct-Session-Id         : "D84FFF000000815556541A"
Sap-Session-Index       : 1

IP Address              : N/A
IP Origin               : N/A
Primary DNS             : N/A
Secondary DNS           : N/A
Primary NBNS            : N/A
Secondary NBNS          : N/A
Address-Pool            : N/A

IPv6 Prefix             : 2001:db8::/64
IPv6 Prefix Origin      : LclPool
IPv6 Prefix Pool        : "pool-v6-1"
IPv6 Del.Pfx.           : N/A
IPv6 Del.Pfx. Origin    : None
IPv6 Del.Pfx. Pool      : ""
IPv6 Address            : N/A
IPv6 Address Origin     : None
IPv6 Address Pool       : ""
Primary IPv6 DNS        : N/A
Secondary IPv6 DNS      : N/A

Radius Session-TO       : N/A
Radius Class            :
Radius User-Name        :
-------------------------------------------------------------------------------
Number of sessions : 1
===============================================================================
```

ICMP6 debugging can be used to show the SLAAC address assignment process.

```
debug
    router
        ip
            icmp6
        exit
    exit

23673 2015/05/14 21:38:38.04 UTC MINOR: DEBUG #2001 TIP
"TIP: ICMP6_PKT
ICMP6 egressing on grp-int-01 (ies1):
   fe80::101 -> ff02::1
   Type: Router Advertisement (134)
```

```
   Code: No Code (0)
      Hop Limit     : 64
      Flags         :
      Retrans Time  : 0
      Def Life Time : 4500
      Reachable Time: 0
      Option  : Prefix         : 2001:db8::/64
                Flags          : On Link Autoconfig
                Valid Life Time: 86400
                Pref  Life Time: 3600

"
```

Conclusion

7750 SR TPSDA offers a variety of address assignment options such as PPP, DHCPv4, DHCPv6, and SLAAC. These options allow service providers to pick the address assignment scheme that best fits their networks. SLAAC address assignment is an essential IPv6 address assignment protocol. Having to assign a static prefix per subscriber host in advance could be a challenge for operators. This chapter provides a complete configuration example of using the local address server to assign prefixes dynamically to IPoE and PPPoE subscriber hosts. This efficient way to assign SLAAC prefixes to subscribers enables operators to achieve a faster time to market for new IPv6 services.

ESMv4: PPPoE Hosts

This chapter describes advanced IPv4 Enhanced Subscriber Management (ESM) PPPoE host configurations.

Topics in this chapter include:

Applicability

This note is applicable for 7750 SR-C12 and SR-7/12 on IOM2 and higher and was tested on release 13.0.R3 and describes support of PPP termination and aggregation (PTA) hosts. L2TP-hosts are out of scope.

This note is related only to the use of IPv4 hosts.

PPPoE hosts are only supported in a Routed CO model (IES or VPRN) using Ethernet SAPs with null, dot1q or QinQ encapsulation.

Summary

The delivery of services to residential customers encompassing voice, video and data is covered by Triple Play Service Delivery Architecture (TPSDA).

In the TPSDA, a subscriber is defined as a collection of hosts pertaining to a single access connection (for example, DSL line) and identified by a subscriber identifier. A subscriber host is an end user terminal within the subscriber home (PC, set-top box, home gateway) that is identified in the network with a unique (IP address/MAC address) tuple for IPoE or (PPPoE session ID; MAC address) tuple for PPPoE.

The following host types are distinguished:

Static hosts

- ip-mac
- ip-only

Dynamic hosts

- ARP-host
- DHCP-host
- PPPoE-host

This chapter provides configuration and troubleshooting commands for PPPoE-hosts and will use a local user database (LUDB) for host authentication and ESM (Enhanced Subscriber Management) string assignments.

The IP information in this note is retrieved from a Local DHCP server.

The authentication, IP information and ESM strings can come all from an LUDB, a RADIUS server, a (local) DHCP server, or any combination of them. These combinations are out of scope.

Knowledge of the TPSDA concept is assumed throughout this document.

Overview

PPPoE Hosts in a Routed CO Environment

The network topology for a Routed CO environment is displayed in Figure 126.

Figure 126 **Routed CO Network Topology**

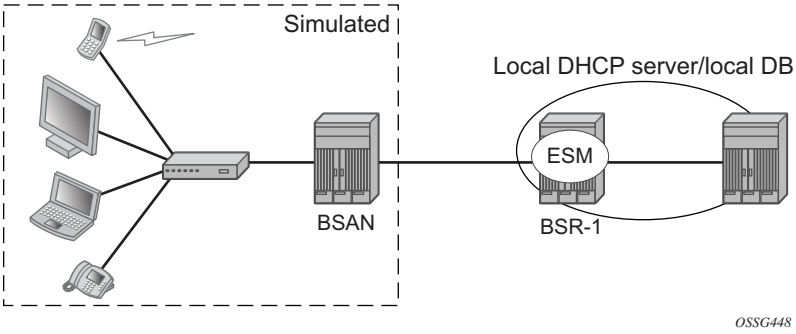

The following configuration tasks should already be configured and are not detailed or explained in this chapter. Refer to the appropriate user guide.

- Basic service router configuration (system interface, IGP, MPLS, BGP)
- Routed CO service topology: VPRN or IES service with subscriber- and group-interface on BSR-1
- ESM
- LUDB (Local User Data Base)
- Local (DHCP) Dynamic Host Configuration Protocol) server

This chapter focuses on PPPoE hosts instantiated in a VPRN service subscriber-interface on BSR-1 (Routed CO).

Note: In case of Routed CO, it is also possible to instantiate the PPPoE hosts in the Base routing instance using an IES service.

Review of the PPPoE Protocol

PPPoE, Point-to-Point Protocol over Ethernet, is a network protocol for encapsulating PPP frames inside Ethernet frames. The protocol is described as an informational RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*, and is based on RFC 1661, *The Point-to-Point Protocol (PPP)*,which provides a standard method for transporting multi-protocol data-grams over point-to-point links.

PPP includes three main components:

- A method for encapsulating multi-protocol datagrams.

- A link control protocol (LCP) for establishing, configuring, and testing the data-link connection.
- A family of network control protocols (NCP) for establishing and configuring different network-layer protocols.

Ethernet networks are packet-based and have no concept of a connection or circuit. By using PPPoE, users can virtually dial from one machine to another over an Ethernet network; establish a point to point connection between them and then transport data packets over the connection.

In a typical wire-line solution with broadband access, PPPoE is used between a client (PC or modem) and a Network Access Server (NAS) (also called Broadband Network Gateway (BNG) or Broadband Service Router (BSR)) through an access node, like a Broadband Service Access Node (BSAN).

PPPoE consists of two phases, the Discovery Stage and the Session Stage.

Discovery Stage

The discovery phase offers a stateless client-server model. When the Discovery Stage completes, both peers know the PPPoE SESSION_ID and the peer's Ethernet address, which together uniquely define the PPPoE session. There are four steps in the Discovery Stage:

Step 1. PPPoE Active Discovery Initiation (PADI)

Initiation (Host broadcast) — This broadcast packet is used by the client to search for an active server (BNG/BSR/NAS) providing access to a service.

Note: Additional attributes on the PADI message could be added if a BSAN is situated between the client and the BRAS.

Step 2. PPPoE Active Discovery Offer (PADO)

Access concentrator unicast — If the access server can provide the service it will 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.

Step 3. PPPoE Active Discovery Request (PADR):

Host unicast — After the client receives a PADO it will send a PADR unicast packet to connect to a server.

Step 4. PPPoE Active Discovery Session-Confirmation (PADS)

Access concentrator unicast — A server will respond to the client with this unicast packet to establish the session and provide the session-id. Once the PADS was provided the Session Stage begins.

Note: Discovery PPPoE Ethernet frames have the ETHER_TYPE field set to the value 0x8863.

PPPoE Tags

IANA has set up a registry of PPPoE tag values (16-bit values). PPPoE tag values already in use are specified as reserved values as shown in Table 16. All other tag values between 0 and 65535 are to be assigned by IANA.

Table 16 **Reserved PPPoE Tags**

Tag Value	Tag Name
0 0x0000	End-Of-List
257 0x0101	Service-Name
258 0x0102	AC-Name
259 0x0103	Host-Uniq
260 0x0104	AC-Cookie
261 0x0105	Vendor-Specific
262 0x0106	Credits
263 0x0107	Metrics
264 0x0108	Sequence Number
272 0x0110	Relay-Session-Id
273 0x0111	HURL
274 0x0112	MOTM
288 0x0120	PPP-Max-Payload
289 0x0121	IP_Route_Add
513 0x0201	Service-Name-Error
514 0x0202	AC-System-Error

Table 16 **Reserved PPPoE Tags (Continued)**

Tag Value	Tag Name
515 0x0203	Generic-Error

Explanations for some PPPoE tags (RFC 2516) are shown in the PPPoE discovery debug messages.

0x0101 Service-Names — This tag indicates that a service name follows. The tag_value is an UTF-8 string that is not null terminated. When the tag_length is zero this tag is used to indicate that any service is acceptable. Examples of the use of the *service-name* tag are to indicate an ISP name or a class or quality of service.

(0x0102) AC-Names — This tag indicates that a string follows which uniquely identifies this particular Access Concentrator unit from all others. It may be a combination of trademark, model, and serial id information, or simply an UTF-8 rendition of the MAC address of the box. It is not null terminated.

(0x0103) Host-Uniq — This tag is used by a host to uniquely associate an access concentrator response (PADO or PADS) to a particular host request (PADI or PADR). The tag_value is binary data of any value and length that the host chooses. It is not interpreted by the Access Concentrator. The host may include a host-uniq tag in a PADI or PADR. If the access concentrator receives this tag, it must include the tag unmodified in the associated PADO or PADS response.

(0x0104) AC-Cookie — This tag is used by the access concentrator to aid in protecting against denial of service attacks. The access concentrator may include this tag in a PADO packet.

If a host receives this tag, it must return the tag unmodified in the following PADR. The tag_value is binary data of any value and length and is not interpreted by the host.

Figure 127 **Discovery Stage Messages**

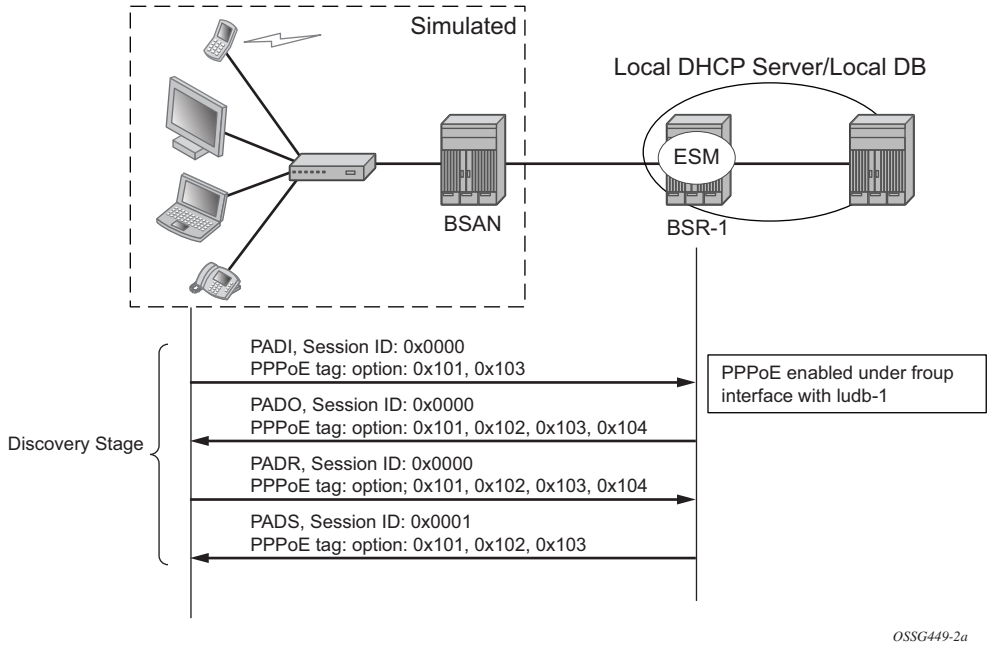

Session Stage

This next stage after Discovery is called the Session Stage. Once the MAC address of the peer is known and a session-id is exchanged, the two end points have all the information needed to start building a point-to-point connection over Ethernet and exchange packets over the connection.

This stage can be divided into to the following sections:

- Setup
- Maintenance
- Termination

Setup

PPP Link Control Protocol (LCP)

Both the NAS and the user open the PPP session based on LCP packets. All post-discovery PPPoE Ethernet frames have the ETHER_TYPE field set to the value 0x8864.

Authentication method and the MRU are negotiated during this phase.

RFC 2516 mandates a maximum negotiated Maximum Receive Unit (MRU) of 1492.

RFC 4638, *Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE),* relaxes this restriction and allows a maximum negotiated MRU greater than 1492 to minimize fragmentation in next-generation broadband networks.

The 7750 implementation follows RFC 4638 when the client implements these extensions.

LCP uses config_request and config_ack/nack to negotiate parameters:

- LCP goes to final state opened when configure-ack is send & received.
- The own options are proposed in configure request.

There are three cases for the LCP negations parameters:

- Peer supports the options and his content.
  - Peer will agree and send config-ack.
- Peer does not support an option
  - Peer will send configure-reject with the option that is not supported.
  - Resend of configure-request without that option.
  - Peer agrees and sends config-ack.
- Peer does support the option but not the content.
  - Peer will send config-nack with the option and his new content.
  - Resend of configure-request with same options but new content.
  - Peer agrees and sends config-ack.

Table 17 **LCP and IPCP Code**

Code	Packet Type
1	Configure-Request
2	Configure-Ack
3	Configure-Nak
4	Configure-Reject
5	Terminate-Request
6	Terminate-Ack
7	Code-Reject
8	Protocol-Reject
9	Echo-Request
10	Echo-Reply
11	Discard-Request
12	Identification
13	Time-Remaining
14	Reset-request CCP
15	Reset-Ack CCP

Figure 128 **LCP Phase Messages**

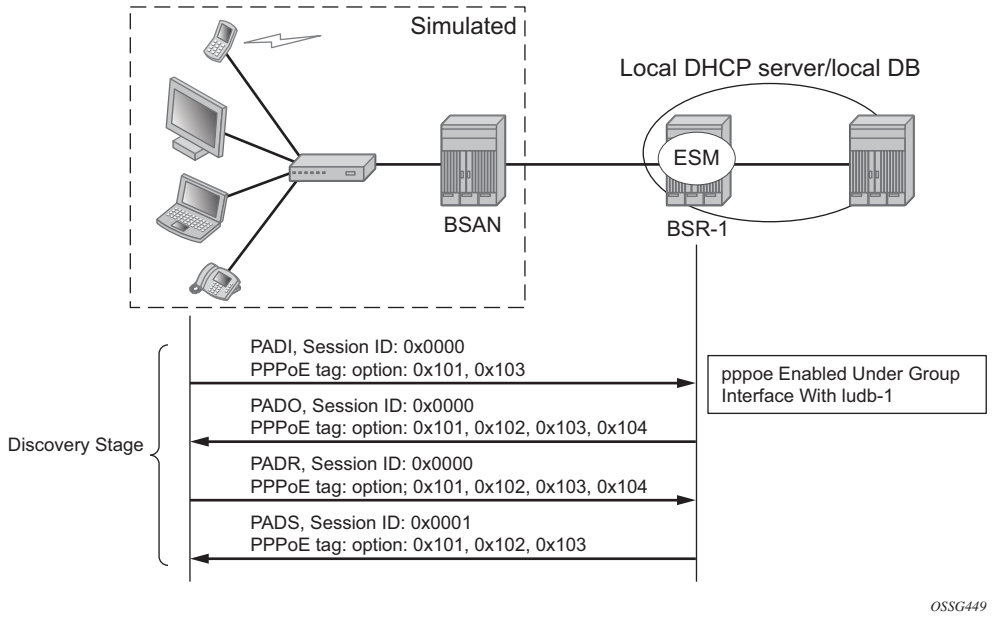

Authentication Phase

The client authenticates itself through PAP (PPP Password Authentication Protocol) or CHAP (Challenge Handshake Authentication Protocol) to check for access permission. For the CHAP authentication, the BSR initiates the authentication as shown in Figure 129.

Note: The password is a hashed output on the link and plain text in a RADIUS Access-Request message.

Figure 129 **CHAP Handshaking Overview Process**

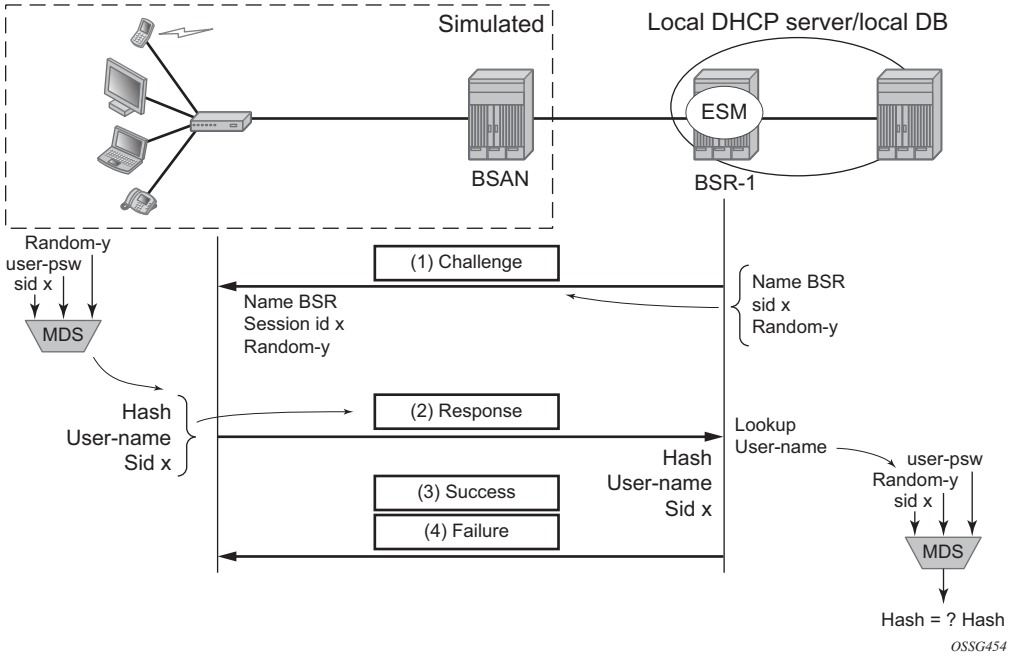

For PAP the client initiates the authentication as shown in Figure 130.

Note: The password is sent as plain text on the link and hashed in a RADIUS Access-Request message.

Figure 130 **PAP Overview Process**

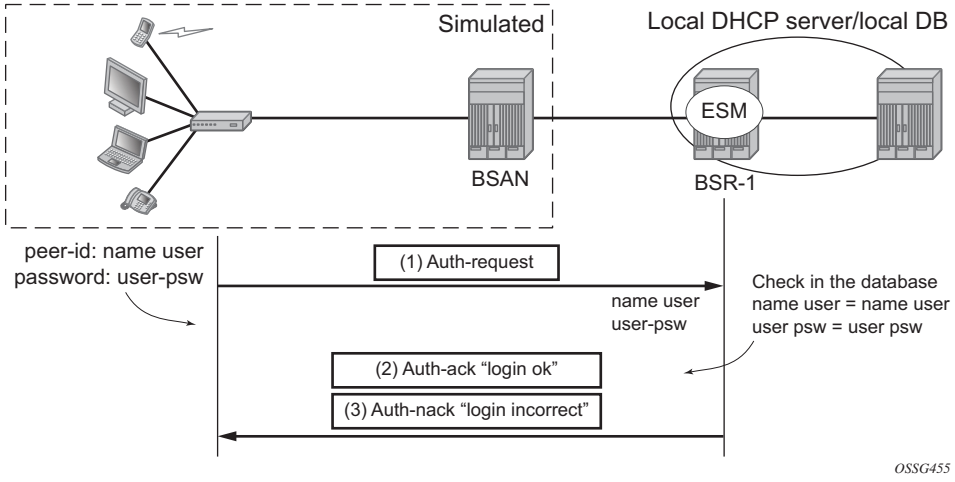

Network-Layer Protocol Phase (PPP IPCP Opening Phase)

At this stage the user requests an IP address to be used for data transmission. During this negotiation the client will also receive a Domain Name Server (DNS), NBNS (Netbios Name Server) address, etc. if they are requested.

Figure 131 **IPCP Phase Messages**

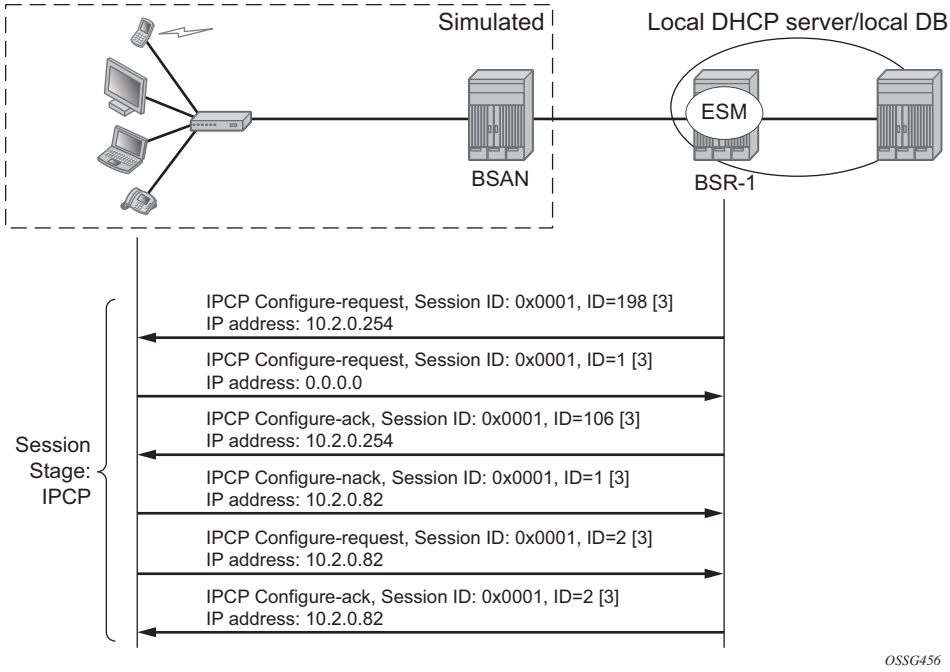

Maintenance

PPP uses keepalives in order to maintain the integrity of the connection. This keepalive mechanism uses an echo-request that is sent to remote PPP peer, following which the remote PPP peer should respond with an echo-reply. The connection is considered down if a number of echo replies are missed. Both sides can initiate keepalives which run independently.

Figure 132 **Keepalive Messages**

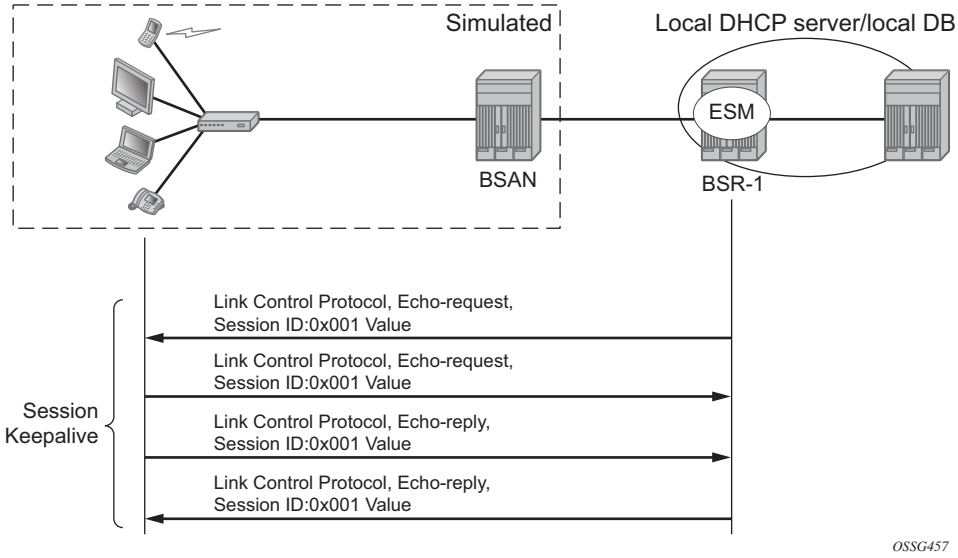

Termination

Link Termination Phase

A PPPoE session can be terminated by either the client or by the BRAS and consists of a Terminate-request followed by a PADT.

Figure 133 **Link Termination Phase**

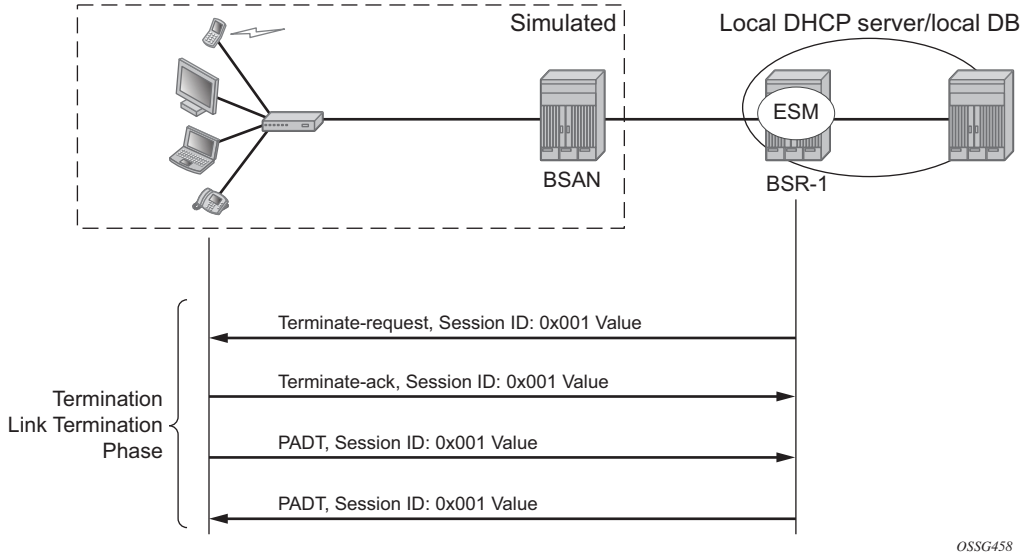

Configuration

PPPoE Host Session: Set-Up, Operation and Release

Enable PPPoE termination under the group-interface context.

Enable the local user database under the PPPoE node of the group-interface.

```
configure
    service
        vprn 1 customer 1 create
            subscriber-interface "sub-int-1" create
                address 10.2.0.254/16
                group-interface "group-int-1" create
                    sap 1/1/1:1 create
                        sub-sla-mgmt
                            sub-ident-policy "sub-id-default"
                            no shutdown
                        exit
                    exit
                    pppoe
                        user-db "ludb-1"
                        no shutdown
                    exit
                exit
            exit
        exit
```

The local user database is configured with the following parameters.

```
configure
    subscriber-mgmt
        local-user-db "ludb-1" create
            ppp
                match-list username
                host "user1" create
                    host-identification
                        username "user1@domain1"
                    exit
                    address pool "pool-1"
                    password chap ALU
                    identification-strings 254 create
                        subscriber-id "PPPoE-host-user1@domain1"
                        sla-profile-string "sla-profile-1"
                        sub-profile-string "sub-profile-1"
                    exit
                    no shutdown
                exit
```

A local DHCP server will be used as a source for the IP addressing of the PPPoE host.

```
configure
    service
        vprn 1
            dhcp
                local-dhcp-server "server-1" create
                    use-pool-from-client
                    pool "pool-1" create
                        subnet 10.2.0.0/16 create
                            exclude-addresses 10.2.0.254 10.2.0.255
                            address-range 10.2.0.1 10.2.0.253
                        exit
                    exit
                    no shutdown
                exit
            exit
```

- The PPPoE policy configuration.

```
*A:BSR-1# configure subscriber-mgmt ppp-policy "ppp-policy-1"
*A:BSR-1>config>subscr-mgmt>ppp-policy# info detail
----------------------------------------------
            no description
            no default-pap-password
            no default-user-name
            disable-cookies
            no force-ppp-mtu-gt-1492
            no ipcp-subnet-negotiation
            keepalive 30 hold-up-multiplier 3
```

```
            no pado-ac-name
            pado-delay 30
            no ppp-initial-delay
            no ppp-mtu
            no lcp-ignore-magic-numbers
            max-sessions-per-mac  63
            mlppp
                no accept-mrru
                no endpoint
                no short-sequence-numbers
            exit
            reply-on-padt
            ppp-authentication pref-chap
            ppp-chap-challenge-length min 32 max 64
            no unique-sid-per-sap
            no re-establish-session
            no reject-disabled-ncp
            no session-timeout
            ppp-options
            exit
-----------------------------------------------
*A:BSR-1>config>subscr-mgmt>ppp-policy#
```

The PPPoE policy defines the parameters which are used in the establishment of the PPPoE session such as:

- Disable-cookies — This parameter disables the use of cookies.
- Keepalive — This command defines the keepalive interval and the number of keepalives that can be missed before the session is declared down for this PPPoE policy.
  - [10 — 300] seconds: Specifies the keepalive interval in seconds.
  - hold-up-multiplier [1 — 5]: Specifies the number of keepalives that can be missed.
- PADO-delay — This parameter configures the delay timeout before sending a PPPoE Active Discovery Offer (PADO) packet.
  - [1 — 30] deciseconds
- PPP-mtu — This parameter configures the maximum PPP MTU size.
  - [512 — 9212]: possible values for MTU size.
- Max-sessions-per-mac — This parameter sets the maximum PPPoE sessions that can be opened for the given MAC address.
  - [1 — 63]: possible PPPoE sessions per MAC address.
- Reply-on-PADT — Some of the PPPoE clients expect reply on PPPoE Active Discovery Terminate (PADT) message before the context of the session is cleared up. To support such client, a command enabling reply to PADT is provided.
  - [Default] **no reply-on-padt**

- PPP-options — This parameter enables the context to configure PPP options which is not supported by default

These parameters will be explained later in details according to its existence in which PPPoE phase.

Note:

- The default policy cannot be modified nor deleted.
- Multiple PPPoE policies may be configured.

- The PPPoE policy is applied within the PPPoE context under the group interface.

```
configure
    service
        vprn 1
            subscriber-interface "sub-int-1" create
                address 10.2.0.254/16
                group-interface "group-int-1" create
                    pppoe
                        user-db "ludb-1"
                        no shutdown
                    exit
                exit
            exit
```

Troubleshooting the PPPoE discovery messages (PADI, PADO, PADR, PADS and PADT) is done with PPPoE debugging:

```
*A:BSR-1# debug service id 1 ppp packet discovery
  - discovery [padi] [pado] [padr] [pads] [padt]
  - no discovery

 <padi>               : keyword - debug PADI packets
 <pado>               : keyword - debug PADO packets
 <padr>               : keyword - debug PADR packets
 <pads>               : keyword - debug PADS packets
 <padt>               : keyword - debug PADT packets

*A:BSR-1#
```

For example:

```
debug
  service
      id 1
          ppp
              packet
                  mode egr-ingr-and-dropped
                  discovery
                  ppp
                  dhcp-client
              exit
          exit
      exit
  exit
exit
```

To display the debugging information, a dedicated log should be created:

```
configure
    log
        log-id 1
            from debug-trace
            to session
            no shutdown
        exit
```

Discovery Stage

The following is an example of PPPoE (PADI discovery packet) debug log output:

```
1 2015/07/01 09:36:09.93 UTC MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: RX Packet
 VPRN 1, SAP 1/1/1:1

 DMAC: ff:ff:ff:ff:ff:ff
 SMAC: 00:00:67:14:01:02
 Ether Type: 0x8863 (Discovery)

 PPPoE Header:
 Version: 1                Type      : 1
 Code   : 0x09 (PADI)      Session-Id: 0x0000 (0)
 Length : 65

 PPPoE Tags:
 [0x0101] Service-Name: "AGILENT"
 [0x0103] Host-Uniq: len = 1, value = 34
 [0x0105] Vendor-Specific: vendor-id = 0x0de9 (ADSL Forum)
    [0x01] Agent-Circuit-Id: "circuit10"
    [0x02] Agent-Remote-Id: "remote10"
    [0x81] Actual-Upstream: 1024
    [0x82] Actual-Downstream: 16384
    [0x90] Access-Loop-Encap: 01 01 00
"
```

PPPoE Policy Parameters

Service name — The client can ask a particular service. Empty means that any service is acceptable. The service name can indicate an ISP name, class, QoS.

```
"PPPoE: RX Packet
   VPRN 1, SAP 1/1/1:1

   DMAC: ff:ff:ff:ff:ff:ff
   SMAC: 00:00:67:14:01:02
   Ether Type: 0x8863 (Discovery)

   PPPoE Header:
   Version: 1                Type      : 1
   Code   : 0x09 (PADI)      Session-Id: 0x0000 (0)
   Length : 55

   PPPoE Tags:
   [0x0101] Service-Name: "AGILENT"
```

The BSR echoes the service name from the PADI message. Empty means that any service is acceptable.

```
2 2015/07/01 09:36:12.90 UTC MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: TX Packet
   VPRN 1, SAP 1/1/1:1

   DMAC: 00:00:67:14:01:02
   SMAC: ea:4b:01:01:00:01
   Ether Type: 0x8863 (Discovery)

   PPPoE Header:
   Version: 1                Type      : 1
   Code   : 0x07 (PADO)      Session-Id: 0x0000 (0)
   Length : 48

   PPPoE Tags:
   [0x0101] Service-Name: "AGILENT"
   [0x0102] AC-Name: "BSR-1"
   [0x0103] Host-Uniq: len = 1, value = 34
"
```

Host-Uniq — The host can include a unique tag of any length inserted In PADI or PADR.The AC should echo back this tag in the PADO or PADS.

```
1 2015/07/01 09:36:09.93 UTC MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: RX Packet
   VPRN 1, SAP 1/1/1:1

   DMAC: ff:ff:ff:ff:ff:ff
   SMAC: 00:00:67:14:01:02
   Ether Type: 0x8863 (Discovery)

   PPPoE Header:
   Version: 1                Type      : 1
```

```
   Code   : 0x09 (PADI)       Session-Id: 0x0000 (0)
   Length : 65

   PPPoE Tags:
   [0x0101] Service-Name: "AGILENT"
   [0x0103] Host-Uniq: len = 1, value = 34
```

The following parameters can optionally be added to the PADI by the PPPoE intermediate agent (BSAN):

Vendor-specific information

- Agent-Circuit-Id
- Agent-Remote-Id
- Access-loop-Encapsulation
- Access loop characteristics (actual-upstream, actual-downstream)

The debug output:

```
  1 2015/07/01 09:36:09.93 UTC MINOR: DEBUG #2001 vprn1 PPPoE
  "PPPoE: RX Packet
   VPRN 1, SAP 1/1/1:1

   DMAC: ff:ff:ff:ff:ff:ff
   SMAC: 00:00:67:14:01:02
   Ether Type: 0x8863 (Discovery)

   PPPoE Header:
   Version: 1                Type      : 1
   Code   : 0x09 (PADI)      Session-Id: 0x0000 (0)
   Length : 55

   PPPoE Tags:
   [0x0101] Service-Name: "AGILENT"
   [0x0103] Host-Uniq: len = 1, value = 34
   [0x0105] Vendor-Specific: vendor-id = 0x0de9 (ADSL Forum)
      [0x01] Agent-Circuit-Id: "circuit10"
      [0x02] Agent-Remote-Id: "remote10"
      [0x81] Actual-Upstream: 1024
      [0x82] Actual-Downstream: 16384
      [0x90] Access-Loop-Encap: 01 01 00
"
```

The cookies can be displayed in the PADO message. This tag of any value and length may be included by the AC and is echoed back by the client to the AC in the next PADR.

```
2 2015/07/01 09:36:12.90 UTC MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: TX Packet
   VPRN 1, SAP 1/1/1:1
```

```
   DMAC: 00:00:67:14:01:02
   SMAC: ea:4b:01:01:00:01
   Ether Type: 0x8863 (Discovery)

   PPPoE Header:
   Version: 1                Type      : 1
   Code   : 0x07 (PADO)      Session-Id: 0x0000 (0)
   Length : 18

   PPPoE Tags:
   [0x0101] Service-Name: "AGILENT"
   [0x0102] AC-Name: "BSR-1"
   [0x0103] Host-Uniq: len = 1, value = 34
   [0x0104] AC-
Cookie: len = 16, value = d7 91 cd b7 3e 51 76 d6 03 0a f2 68 8c da ba 74
"
```

AC-name — Identifies the string that uniquely identifies the access concentrator (AC).

```
2 2015/07/01 09:36:12.90 UTC MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: TX Packet
   VPRN 1, SAP 1/1/1:1

   DMAC: 00:00:67:14:01:02
   SMAC: ea:4b:01:01:00:01
   Ether Type: 0x8863 (Discovery)

   PPPoE Header:
   Version: 1                Type      : 1
   Code   : 0x07 (PADO)      Session-Id: 0x0000 (0)
   Length : 48

   PPPoE Tags:
   [0x0101] Service-Name: "AGILENT"
   [0x0102] AC-Name: "BSR-1"
   [0x0103] Host-Uniq: len = 1, value = 34
"
```

When **disable-cookies** is configured, the use of cookies will be disabled, when omitted the **no-disable-cookies** will be used.

```
*A:BSR-1# configure subscriber-mgmt ppp-policy "ppp-policy-1" disable-cookies
```

The cookies are encoded back by the client in the next PADR message.

PPPoE hosts are authenticated based on username-password information (PAP/CHAP authentication) or on information embedded in the PADI message PADI authentication).

The **min** and **max** values for the **ppp-chap-challenge** are defined when enabling **ppp-chap-challenge length**. When omitted, a **min** of 32 and **max** of 64 are used.

```
*A:BSR-1# configure subscriber-mgmt ppp-policy "ppp-policy-1" ppp-chap-challenge-
length
  - ppp-chap-challenge-length min <minimum-length> max <maximum-length>
  - no ppp-chap-challenge-length

 <minimum-length>     : [8..64]
 <maximum-length>     : [8..64]


*A:BSR-1#
```

To complete the discovery phase, the server must provide a session-id to the client and the SR OS allocates session-id 1 for different MACs.

Note: In the VLAN per service model (N: 1 VLAN) where the MACs are the same and the PPPoE interworking will be done at the BSAN.

```
*A:BSR-1# show service id 1 pppoe session
===============================================================================
PPPoE sessions for svc-id 1
===============================================================================
Sap Id               Mac Address       Sid    Up Time         Type
    IP/L2TP-Id/Interface-Id                                      MC-Stdby
-------------------------------------------------------------------------------
1/1/1:1              00:00:67:14:01:02 1      0d 00:01:32     local
    10.2.0.1
-------------------------------------------------------------------------------
Number of sessions   : 1
===============================================================================
*A:BSR-1#
```

The 7750 has the possibility to delay the sending of the PADO message to the client. This feature could be used if the client is dual homed to two BSRs and is explained later in the document.

When PADO-delay is configured, the configured value equals the delay timeout before sending PADO, when omitted the PADO-delay value of 0 msec will be used.

```
*A:BSR-1# configure subscriber-mgmt ppp-policy "ppp-policy-1" pado-delay
  - no pado-delay
  - pado-delay <deci-seconds>

 <deci-seconds>       : [1..30]


*A:BSR-1#
```

Session Stage

LCP

During the link establishment phase client and server negotiate options and need to come to an agreement on these options. Options that are unknown by the peer are rejected whereas known options with unknown content are nack'd. In the later case the peer needs to resend the same option but with another content. In case of a reject the peer should remove that option. An Ack will be send if there is a full agreement.

One of the more important options that is exchanged is the maximum receive unit (MRU) and the authentication protocol that will be used later in the authentication phase. The first option, the MRU value (minus overhead) is sent from the BSR towards the client and is the lowest value between the port MTU and the optional configured ppp-mtu in the ppp-policy.

RFC 2516 mandates a maximum negotiated Maximum Receive Unit (MRU) of 1492 but RFC 4638 relaxes this restriction and allows a maximum negotiated MRU greater than 1492 to minimize fragmentation in next-generation broadband networks. The 7750 SR and 7710 SR implementation follows RFC 4638 when the client implements these extensions.

If a PPPoE client wants to use MRU>1492 in the LCP-config request it should include the **ppp-max-payload** tag with the higher MTU value in the initial PADI message.

```
*A:BSR-1# configure subscriber-mgmt ppp-policy "ppp-policy-1" ppp-mtu
  - no ppp-mtu
  - ppp-mtu <mtu-bytes>

 <mtu-bytes>          : [512..9212]

*A:BSR-1#
```

PPPoE debug output.

```
5 2015/07/01 09:36:12.91 UTC MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: TX Packet
   VPRN 1, SAP 1/1/1:1

   DMAC: 00:00:67:14:01:02
   SMAC: ea:4b:01:01:00:01
   Ether Type: 0x8864 (Session)

   PPPoE Header:
   Version: 1                Type      : 1
   Code   : 0x00             Session-Id: 0x0001 (1)
   Length : 21
```

```
   PPP:
   Protocol  : 0xc021 (LCP)
   Code      : 1 (Configure-Request)
   Identifier: 216            Length    : 19

   Options:
   [1] MRU: 1492
   ---snipped ---
```

The second important option, the authentication method used in the authentication phase is exchanged between client and server and can be PAP or CHAP authentication. The authentication method is not exchanged when PADI authentication is done. PADI authentication means that the BSR will authenticate the user based on parameters in the PADI message. Authentication based on PADI and PAP/CHAP is possible.

Debug output example for CHAP authentication protocol.

```
8 2015/07/01 09:36:12.91 UTC MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: RX Packet
   VPRN 1, SAP 1/1/1:1

   DMAC: ea:4b:01:01:00:01
   SMAC: 00:00:67:14:01:02
   Ether Type: 0x8864 (Session)

   PPPoE Header:
   Version: 1                 Type      : 1
   Code   : 0x00              Session-Id: 0x0001 (1)
   Length : 21

   PPP:
   Protocol  : 0xc021 (LCP)
   Code      : 2 (Configure-Ack)
   Identifier: 216            Length    : 19

   Options:
   [1] MRU: 1492
   [3] Authentication-Protocol: 0xc223 (CHAP), Algorithm = 5 (MD5)
   [5] Magic-Number: 0x7f68ddea
"
```

Debug output example for PAP authentication protocol.

```
33 2015/07/01 09:40:46.21 UTC MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: RX Packet
   VPRN 1, SAP 1/1/1:1

   DMAC: ea:4b:01:01:00:01
   SMAC: 00:00:67:13:01:02
   Ether Type: 0x8864 (Session)

   PPPoE Header:
   Version: 1                Type      : 1
   Code   : 0x00             Session-Id: 0x0001 (1)
   Length : 20

   PPP:
   Protocol  : 0xc021 (LCP)
   Code      : 2 (Configure-Ack)
   Identifier: 74            Length    : 18

   Options:
   [1] MRU: 1492
   [3] Authentication-Protocol: 0xc023 (PAP)
   [5] Magic-Number: 0x79737c8b
"
```

- Fallback case chap->pap

For user authentication, with pap-chap-access, always try CHAP first; if that doesn't succeed, try PAP.

The option to be used first (CHAP/PAP) is defined when enabling the ppp-authentication. When omitted, CHAP is preferred always.

```
A:BSR-1# configure subscriber-mgmt ppp-policy "ppp-policy-1" ppp-authentication
  - no ppp-authentication
  - ppp-authentication {pap|chap|pref-chap|pref-pap}
```

PPPoE clients that implement undocumented options also require an agreement on those unknown options. By default, the 7750 SR will reject unknown options but the **ppp-option** feature in the **pppoe-policy** allows for support of undocumented client LCP or IPCP/IPv6CP options. If the received LCP or IPCP/IPv6CP option matches the configured options in the pppoe-policy an ack will be send instead of a reject.

```
A:BSR-1# configure subscriber-mgmt ppp-policy "ppp-policy-1" ppp-options custom-
option
  - 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>

 <protocol>           : lcp|ipcp|ipv6cp
 <option-number>      : [0..255]
```

```
 <ip-address>          : a.b.c.d
 <ascii-string>        : [127 chars max]
 <hex-string>          : [0x0..0xFFFFFF...(max 254 hex nibbles)]

A:BSR-1#
```

Troubleshooting the PPPoE LCP session messages is done with PPPoE debugging:

```
A:BSR-1# debug service id 1 ppp packet ppp
  - no ppp
  - ppp [lcp] [pap] [chap] [ipcp] [ipv6cp]

 <lcp>                 : keyword - debug LCP packets
 <pap>                 : keyword - debug PAP packets
 <chap>                : keyword - debug CHAP packets
 <ipcp>                : keyword - debug IPCP packets
 <ipv6cp>              : keyword - debug IPv6CP packets

A:BSR-1#
```

At the end of LCP, authentication is started.

The 7750/7710 BSR supports three main methods for PPPoE authentication.

- PADI or PAP/CHAP authentication through RADIUS
- PADI or PAP/CHAP authentication through a LUDB
- PADI authentication via LUDB then PAP/CHAP pre-authentication through RADIUS

DHCP server authentication will not be explained in this section because this is more authorization than authentication

The flow chart of the PPPoE host authentication process is shown in Figure 134 and starts with the request to the middle left.

Figure 134 **Authentication Flow Chart**

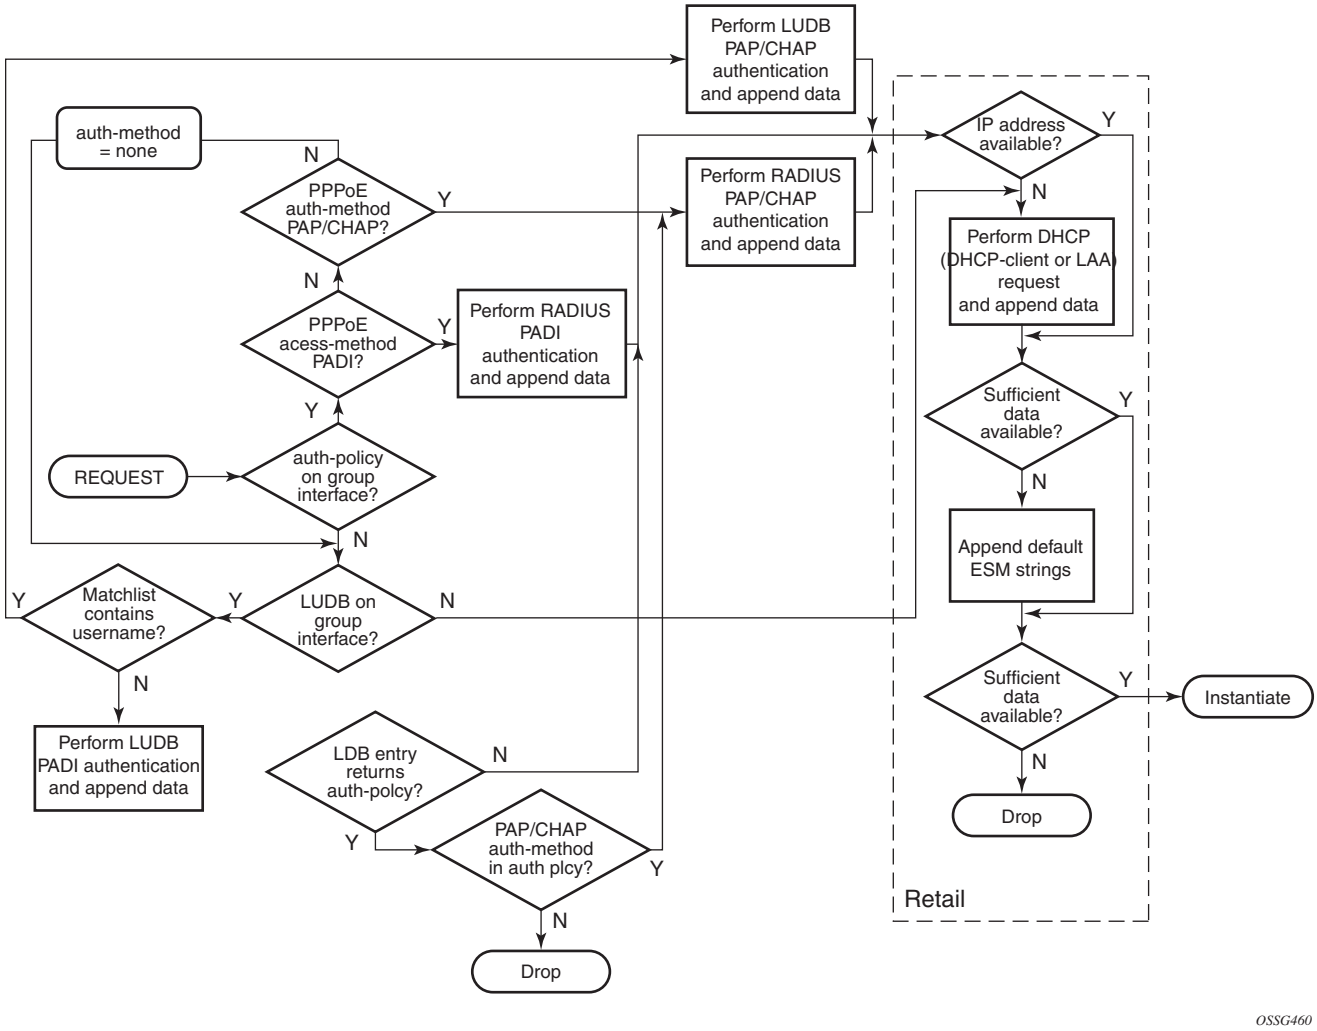

- PPPoE users get authenticated via the LUDB if this LUDB is configured under the group-interface.

```
configure
    service
        vprn 1
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    pppoe
                        user-db "ludb-1"
```

- Users get authenticated via RADIUS if we have an authentication-policy under the same group-interface. RADIUS has precedence if both are configured.

```
configure
    service
        vprn 1
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    authentication-policy "auth-1"
                    pppoe
                        user-db "ludb-1"
                        no shutdown
                    exit
                exit
```

- Users that get authenticated via the LUDB can still go to RADIUS if we move the authentication-policy from the group-interface to the LUDB.

```
configure
 subscriber-mgmt
        local-user-db "ludb-1" create
            ppp
                match-list username
                host "user1" create
                    auth-policy "auth-1"
```

This last mechanism could be used to pick up parameters like **pado-delay** or to check some variables such as **circuit-id**, **remote-id** from the LUDB during discovery phase and subsequently to use RADIUS for PAP/CHAP authentication.

```
*A:BSR-1# configure subscriber-mgmt local-user-db "ludb-1" ppp host "user2" host-
identification
  - host-identification

 [no] circuit-id      - Configure the circuit id of this host
 [no] encap-tag-range - Configure the SAP encap range tags for this host
 [no] mac             - Configure the MAC address of this host
 [no] remote-id       - Configure the remote id of this host
 [no] sap-id          - Configure the SAP identifier of this host
 [no] service-name    - Configure the service name of this host
 [no] username        - Configure the user name of this host

*A:BSR-1#
```

- Both RADIUS and LUDB support PADI or PAP/CHAP authentication.
- PPPoE users that are authenticated through the LUDB and have in the LUDB a match-list other than **username** will get authenticated based on PADI parameters like mac, circuit-id, remote-id.

```
*A:BSR-1# configure subscriber-mgmt local-user-db "ludb-1" ppp match-list
  - no match-list
  - match-list <ppp-match-type-1> [<ppp-match-type-2>...(up to 3 max)]
```

```
 <ppp-match-type>     : circuit-id|mac|remote-id|sap-id|encap-tag-range|service-
name|username

*A:BSR-1#
```

- PPPoE users that have in the LUDB a match-list equal to **username** will use the PAP/CHAP authentication method.

```
configure
    subscriber-mgmt
        local-user-db "ludb-1" create
            ppp
                match-list username
                host "user1" create
                    host-identification
                        username "user1@domain1"
                    exit
                    address pool "pool-1"
                    password chap ALU
                    identification-strings 254 create
                        subscriber-id "PPPoE-host-user1@domain1"
                        sla-profile-string "sla-profile-1"
                        sub-profile-string "sub-profile-1"
                    exit
                    no shutdown
                exit
```

- PPPoE users that are authenticated through RADIUS and have in the authentication policy, a pppoe-access-method equal to PADI will use the **mac** or **circuit-id** information from the PADI in their request to RADIUS.

```
*A:BSR-1# configure subscriber-mgmt authentication-policy "auth-1" pppoe-access-
method
  - no pppoe-access-method
  - pppoe-access-method {none|padi|pap-chap}

*A:BSR-1#
```

- The selection for mac or circuit-id can be altered via the parameter user-name-format.

```
*A:BSR-1# configure subscriber-mgmt authentication-policy "auth-1" user-name-format
  - no user-name-format
  - 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>]
```

```
 <format>             : ascii-converted-circuit-id|ascii-converted-tuple|circuit-
id|dhcp-client-vendor-opts|mac|mac-giaddr|ppp-user-name|tuple
 <domain-name>        : max 128 chars, no @ needed
 <mac-format>         : (only when format is dhcp-client-vendor-opts)
                        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
```

PPPoE users that are authenticated through RADIUS and have in the authentication policy a pppoe-access-method equal to pap-chap use the username from the authentication phase in their request to RADIUS. The parameter user-name-format is irrelevant in this last case.

RADIUS Authentication

When authentication is provided through RADIUS, two methods can be used to authenticate the PPPoE session.

- PADI authentication
- PAP/CHAP authentication

The RADIUS policy specifies which parameters are provided in the RADIUS access-request message.

The following parameters can be configured:

- Circuit-id: Provided through the PADI/PADR PPPoE relay vendor specific tag as specified in TR-101 of the DSL Forum.
- Remote-id: Provided through the PADI/PADR PPPoE relay vendor specific tag as specified in TR-101 of the DSL Forum.
- NAS-port-id: SAP ID on which the PPPoE session terminates (e.g. 1/1/3:1).
- NAS-identifier: System name of the NAS or BNG.
- PPPoE-service-name: Provided through the PPPoE PADI packet.
- access-loop-options: Provided through the PAD/PADR PPPoE extensions as specified in TR-101 of the DSL Forum.

The option to use PADI authentication or PAP/CHAP authentication is selected with the following configuration in the RADIUS policy:

```
*A:BSR-1# configure subscriber-mgmt authentication-policy "auth-1" pppoe-access-
method
  - no pppoe-access-method
  - pppoe-access-method {none|padi|pap-chap}
```

```
*A:BSR-1#
```

When PADI authentication is used for PPPoE termination the MAC address or the PPPoE relay tag (Circuit-ID) or a combination of MAC address and circuit ID can be used to identify the subscriber in the RADIUS server.

To determine the information to use in the RADIUS Access-Request message is configured in the RADIUS policy using the following command:

```
*A:BSR-1# configure subscriber-mgmt authentication-policy "auth-1" user-name-format
  - no user-name-format
  - 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>]

 <format>             : ascii-converted-circuit-id|ascii-converted-tuple|circuit-
id|dhcp-client-vendor-opts|mac|mac-giaddr|ppp-user-name|tuple
 <domain-name>        : max 128 chars, no @ needed
 <mac-format>         : (only when format is dhcp-client-vendor-opts)
                        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
*A:BSR-1#
```

Local User Database Authentication

A second authentication option for PPPoE termination is to use the local user database of the BSR 7750 for PAP/CHAP authentication (this option will be used in this note). With this authentication method the client's PPPoE session is authenticated locally on the BSR without any constraint of an external radius server.

The local user database is configured with the following parameters:

```
configure
    subscriber-mgmt
        local-user-db "ludb-1" create
            ppp
                match-list username
                host "user1" create
                    host-identification
                        username "user1@domain1"
                    exit
                    address pool "pool-1"
                    password chap ALU
                    identification-strings 254 create
                        subscriber-id "PPPoE-host-user1@domain1"
                        sla-profile-string "sla-profile-1"
```

```
                    sub-profile-string "sub-profile-1"
                exit
                no shutdown
            exit
```

With this authentication method, authentication can be provided by the username/password.

To enable the local authentication method, the local user database is configured under the PPPoE node of the group-interface as shown below.

```
configure
    service
        vprn 1
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    pppoe
                        policy "ppp-policy-1"
                        user-db "ludb-1"
                        no shutdown
                    exit
                exit
            exit
```

Refer to the LUBD Basics chapter for further information.

DHCP Client Authentication

A third authentication method for PPPoE termination is to perform PPPoE to DHCP transformation (where the 7750 SR acts as a DHCP client on behalf of the PPP session) and to use a DHCP server for session authentication. This method is useful when a similar authentication is used for DHCP based clients.

The PPPoE to DHCP authentication method can provide authentication on the basis of MAC address, circuit ID or remote ID.

IPCP

IP and DNS information can be obtained from different sources like LUDB and RADIUS for fixed IP addressing or (local) DHCP for dynamic IP pool management.

If IP information is returned from a DHCP server, then the PPPoE options such as the DNS name are retrieved from the DHCP ACK and provided to the PPPoE client.

Local DHCP Server

In this note, a local DHCP server will be used as a source for the IP addressing of the PPPoE host.

```
configure
  service
      vprn 1
          dhcp
              local-dhcp-server "server-1" create
                  use-gi-address
                  pool "pool-1" create
                      subnet 10.2.0.0/16 create
                          exclude-addresses 10.2.0.254 10.2.0.255
                          address-range 10.2.0.1 10.2.0.253
                      exit
                  exit
                  no shutdown
              exit
          exit
```

To check the DHCP server summary:

```
*A:BSR-1# show router 1 dhcp local-dhcp-server "server-1" summary
===============================================================================
DHCP server server-1  router 1
===============================================================================
Admin State            : inService
Operational State      : inService
Persistency State      : shutdown
--- snipped ---
-------------------------------------------------------------------------------
Subnet                 Free     %    Stable   Declined Offered  Rem-pend Drain
-------------------------------------------------------------------------------
10.2.0.0/16            252      99%  1        0        0        0        N

Totals for pool        252      99%  1        0        0        0
-------------------------------------------------------------------------------

Totals for server      252      99%  1        0        0        0

-------------------------------------------------------------------------------
Interface associations
Interface                       Admin
-------------------------------------------------------------------------------
local-dhcp-server-1             Up
--- snipped ---
===============================================================================
*A:BSR-1#
```

To debug the DHCP server:

```
debug
    router "1"
        ip
            dhcp
                detail-level medium
                mode ingr-and-dropped
            exit
        exit
    exit
```

Keepalive

The keepalive timer (defined in seconds) and the hold-up-multiplier are defined when enabling keepalives. When omitted, a 30 second keepalive timer and three hold-up-multipliers are used.

```
*A:BSR-1# configure subscriber-mgmt ppp-policy "ppp-policy-1"  keepalive
  - keepalive <seconds> [hold-up-multiplier <multiplier>]
  - no keepalive

 <seconds>            : [4..300]
 <multiplier>         : [1..5]

*A:BSR-1#

[4-300] seconds: interval between LCP Echo Requests
hold-up-multiplier [1-
5] : Number of missed replies before the PPPoE session is consid- ered dead.
```

To check the keepalive statistics:

```
*A:BSR-1# show service id 1 pppoe session session-
id 1 mac 00:00:67:14:01:02 statistics
===============================================================================
PPPoE sessions for svc-id 1
===============================================================================
Sap Id              Mac Address       Sid    Up Time         Type
    IP/L2TP-Id/Interface-Id                                    MC-Stdby
-------------------------------------------------------------------------------
1/1/1:1             00:00:67:14:01:02 1      0d 00:18:51     local
    10.2.0.1

Packet Type               Received        Transmitted
-------------------------------------------------------------------------------
--- snipped ---
LCP Echo-Request          113             0
LCP Echo-Reply            0               113
--- snipped ---

Number of sessions   : 1
```

```
===============================================================================
*A:BSR-1#
```

The 7750 BSR supports an optimized implementation of keepalive mechanism; this is a mechanism where client and/or server can check the aliveness of the peer. This LCP echo-request is sent on expiration of a timer, derived from the configured **pppoe-policy keepalive** value.

An LCP echo reply is returned to the client after a LCP echo request is received and the above timer on the BSR is reset to the initial keepalive value.

The above mechanism results in an optimized mechanism if the keepalive timers from the client are smaller than the configured values on the BSR.

The client or BSR will terminate the session with a PADT if no LCP echo-reply is received within the time specified by the hold-up-multiplier.

Example for Echo Request from BSR and Echo Reply from the PPPoE host.

```
Ethernet II, Src: TimetraN_90:f8:6a (00:03:fa:90:f8:6a), Dst: Soft*Rit_14:01:02 (00:
00:67:14:01:02)
802.1Q Virtual LAN, PRI: 5, CFI: 0, ID: 1
PPP-over-Ethernet Session
Point-to-Point Protocol
PPP Link Control Protocol
    Code: Echo Request (0x09)
    Identifier: 0x01
    Length: 8
    Magic number: 0x2e538af6

No.     Time        Source                Destination           Protocol Info
      4 30.000280   TimetraN_90:f8:6a     Soft*Rit_14:01:02     PPP LCP  Echo Reply

Ethernet II, Src: TimetraN_90:f8:6a (00:03:fa:90:f8:6a), Dst: Soft*Rit_14:01:02 (00:
00:67:14:01:02)
802.1Q Virtual LAN, PRI: 5, CFI: 0, ID: 1
PPP-over-Ethernet Session
Point-to-Point Protocol
PPP Link Control Protocol
    Code: Echo Reply (0x0a)
    Identifier: 0x0b
    Length: 8
    Magic number: 0x2e538af6
```

To check the PPPoE session for a particular service, use the **show service id <***service-id***> pppoe session** command. Detailed output as well as additional output filtering is available:

```
*A:BSR-1# show service id 1 pppoe session
  - session [interface <ip-int-name|ip-address> | sap <sap-id>] [type <pppoe-
session-type>] [session-id <session-id>] [mac <ieee-address>] [ip-address
    <ip-prefix[/prefix-length]>] [port <port-id>] [no-inter-dest-id | inter-dest-
```

```
id <intermediate-destination-id>] [detail|statistics]
  - session l2tp-connection-id <connection-id> [detail|statistics]

*A:BSR-1# show service id 1 pppoe session detail

===============================================================================
PPPoE sessions for svc-id 1
===============================================================================
Sap Id               Mac Address       Sid    Up Time         Type
    IP/L2TP-Id/Interface-Id                                      MC-Stdby
-------------------------------------------------------------------------------
1/1/1:1              00:00:67:14:01:02 1      0d 00:19:34     local
    10.2.0.1


LCP State            : Opened
IPCP State           : Opened
IPv6CP State         : Closed
PPP MTU              : 1492
PPP Auth-Protocol    : CHAP
PPP User-Name        : user1@domain1

Subscriber-interface : sub-int-1
Group-interface      : group-int-1

IP Origin            : dhcp
DNS Origin           : none
NBNS Origin          : none

Subscriber           : "PPPoE-host-user1@domain1"
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      : "EA4BFF000000005593B48C"
Sap-Session-Index    : 1

IP Address           : 10.2.0.1/32
Primary DNS          : N/A
Secondary DNS        : N/A
Primary NBNS         : N/A
Secondary NBNS       : N/A
Address-Pool         : pool-1

--- snipped ---

Ignoring DF bit      : false

Circuit-Id           : circuit10
Remote-Id            : remote10

Radius Session-TO    : N/A
Radius Class         :
Radius User-Name     : user1@domain1
Logical-Line-Id      :
Service-Name         :
Data link            : ethernet
```

```
Encaps 1             : untagged-ethernet
Encaps 2             : not-available
Origin               : tags
Link Rate Down       : 16384
Rate Origin          : tags
-------------------------------------------------------------------------------
Number of sessions   : 1
===============================================================================
*A:BSR-1#
```

An event will be generated when a PPPoE host has been created in the system.

```
46 2015/07/01 09:36:12.92 UTC WARNING: SVCMGR #2500 Base Subscriber created
"Subscriber PPPoE-host-user1@domain1 has been created in the system"
```

The PPPoE host will appear in the subscriber-host table for the service with origin set to IPCP.

```
*A:BSR-1# show service id 1 subscriber-hosts
===============================================================
Subscriber Host table
===============================================================
Sap                   Subscriber
  IP Address
    MAC Address          PPPoE-SID Origin       Fwding State
---------------------------------------------------------------
1/1/1:1               PPPoE-host-user1@domain1
  10.2.0.1
    00:00:67:14:01:02    1         IPCP         Fwding
---------------------------------------------------------------
Number of subscriber hosts : 1
===============================================================
*A:BSR-1#
```

A host route (/32) for its IP address is inserted in the routing table towards the appropriate group-interface.

```
*A:BSR-1# show router 1 route-table
===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.2.0.0/16                                   Local   Local     00h24m10s  0
       sub-int-1                                                  0
10.2.0.1/32                                   Remote  Sub Mgmt  00h20m35s  0
       [group-int-1]                                              0
172.16.0.1/32                                 Local   Local     00h24m27s  0
       local-dhcp-server-1                                        0
-------------------------------------------------------------------------------
No. of Routes: 3
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
```

```
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
*A:BSR-1#
```

To advertise the PPPoE host subnets to other protocol/network, a policy statement should be defined with using **from protocol direct**.

```
configure router policy-options
        begin
            policy-statement "policy-1"
                entry 10
                    from
                        protocol direct
                    exit
                    --- snipped ---
```

Terminate

Some PPPoE clients expect a reply on PADT message before the context of the session is cleared up. To support such clients, a command enabling reply to PADT is configured.

When reply-on-padt is configured, the BSR will reply with PADT message, when omitted no PADT will be sent from the BSR as a reply on the client's PADT.

```
*A:BSR-1# configure subscriber-mgmt ppp-policy "ppp-policy-1" reply-on-padt
```

The PPPoE debug output:

```
77 2010/01/07 20:31:50.32 UTC MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: TX Packet
   VPRN 1, SAP 1/1/3:1

   DMAC: 00:00:67:13:01:02
   SMAC: 00:03:fa:90:f8:6a
   Ether Type: 0x8863 (Discovery)

   PPPoE Header:
   Version: 1                Type      : 1
   Code   : 0xa7 (PADT)      Session-Id: 0x0001 (1)
   Lengt
```

A PPPoE host can be manually deleted from the system using following clear command:

```
*A:BSR-1# clear service id 1 ppp session
  - session [sap-id <sap-id>] [interface <ip-int-name|ip-address>] [mac <ieee-
address>] [session-id <session-id>] [type
    <pppoe-session-type>] [ip-address <ip-prefix[/prefix-length]>] [port <port-
id>] [inter-dest-id
    <intermediate-destination-id>] [no-inter-dest-id] [user-name <user-name>] [sub-
ppp-type {oa|oe|oeoa|ol2tp}] [no-padt]
  - session all [no-padt]

*A:BSR-1# clear service id 1 ppp session sap-id 1/1/1:1 mac 00:00:67:14:01:02
```

The causes for terminating the PPP session can be:

- **Admin Reset** — Use the **clear** command or a RADIUS Disconnect Request.

  `TERMINATE CAUSE [49] 4 Admin Reset(6)`
- **User Request** — User disconnects the session.

  `TERMINATE CAUSE [49] 4 User Request(1)`
- **Accounting OFF**
  - When accounting policy has been removed from sap/interface/sub-profile.
  - vprn service which is transporting accounting information has been shutdown.
  - The last RADIUS accounting server has been removed from already applied accounting policy.

  `TERMINATE CAUSE [49] 4 NAS Request(10)`
- **PPPoE keepalive timeout**

  `TERMINATE CAUSE [49] 4 Lost Carrier(2)`
- **RADIUS session timeout**

PPPoE Hosts Advanced Topics

QoS Aspects

VLAN based downstream PPPoE control traffic is generated by default with dot1p value 7 .This value can be overruled with the following commands:

In case of the PPPoE hosts instantiated in the Base routing instance using an IES service.

```
*A:BSR-1# configure router sgt-qos application pppoe dot1p 5
```

In case of the PPPoE hosts instantiated in a VPRN service subscriber-interface.

```
*A:BSR-1# configure service vprn 1 sgt-qos application pppoe dot1p 5
```

The **show router sgt-qos** command displays the configured and default DSCP and default dot1p values per application. Since PPPoE is a Layer 2 protocol we will see only the dot1p settings. The default dot1p value **none** corresponds with value 7.

```
*A:BSR-1# show router 1 sgt-qos application pppoe
===============================================================================
Dot1p Application Values
===============================================================================
Application         Dot1p Value                  Default Dot1p Value
-------------------------------------------------------------------------------
pppoe               5                            none
===============================================================================
*A:BSR-1#
```

Limiting the Number of PPPoE Hosts

The maximum number of PPPoE sessions can be controlled by the parameters session-limit, sap-session-limit, host-limit, multi-sub-sap limit and max-sessions-per mac.

session-limit — The maximum number of PPPoE sessions for an IES/VPRN group-interface is defined when enabling session-limit. When omitted, a single PPPoE session is allowed.

```
configure
    service
        vprn 1
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    pppoe
                        session-limit 1
```

A trap is generated when trying to instantiate a new PPPoE session while the configured number of sessions is reached.

Note: The discovery phase is completed before the check on the session-limit is performed.

```
47 2015/07/02 09:17:11.21 UTC WARNING: PPPOE #2001 vprn1 PPPoE session failure
"PPPoE session failure on SAP 1/1/1:1 in service 1 -
 Reached the interface session limit (1) for "group-int-1""
```

PPPoE debug output:

```
23 2015/07/02 09:17:11.21 UTC MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: Dropped Packet
   VPRN 1, SAP 1/1/1:1

   Problem: Reached the interface session limit (1) for "group-int-1"

   DMAC: ff:ff:ff:ff:ff:ff
   SMAC: 00:00:67:13:01:02
   Ether Type: 0x8863 (Discovery)
```

sap-session-limit

The maximum number of PPPoE sessions per SAP for an IES/VPRN group-interface is defined when enabling sap-session-limit. When omitted, a single PPPoE session per SAP is allowed:

```
configure
    service
        vprn 1
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    pppoe
                        sap-session-limit 1
```

A trap is generated when trying to instantiate a new PPPoE session while the configured number of the sessions per sap is reached.

```
48 2015/07/02 09:22:26.99 UTC WARNING: PPPOE #2001 vprn1 PPPoE session failure
"PPPoE session failure on SAP 1/1/1:1 in service 1 - Reached the per-
SAP session limit (1) for "group-int-1""

PPPoE debug output:


25 2015/07/02 09:22:26.99 UTC MINOR: DEBUG #2001 vprn1 PPPoE
"PPPoE: Dropped Packet
   VPRN 1, SAP 1/1/1:1

   Problem: Reached the per-SAP session limit (1) for "group-int-1"

   DMAC: ff:ff:ff:ff:ff:ff
   SMAC: 00:00:67:13:01:02
   Ether Type: 0x8863 (Discovery)
```

max-sessions-per-mac

The BSR 7750 implementation defines a unique PPPoE session based on the PPPoE SESSION_ID and the client's MAC address.

The maximum number of PPPoE sessions per mac is defined when enabling max-sessions-per-mac. When omitted, a single PPPoE session per mac is allowed.

```
configure
    subscriber-mgmt
        ppp-policy "ppp-policy-1"
               max-sessions-per-mac 63
```

Although the command is max-session-per-mac, actually it means the maximum number of supported sessions-per-MAC-per-SAP especially in N: 1 VLAN model.

A trap is generated when trying to instantiate a new PPPoE session per the same mac while the configured number of max-sessions-per-mac is reached.

```
52 2015/07/02 09:32:11.41 UTC WARNING: PPPOE #2001 vprn1 PPPoE session failure
"PPPoE session failure on SAP 1/1/1:1 in service 1 -
 Reached the maximum number (1) of PPPoE sessions for MAC 00:00:67:14:01:02"
```

Host-limit

The maximum number of PPPoE hosts is defined when enabling host-limit. When omitted, a single host is allowed.

```
configure
    subscriber-mgmt
        sla-profile "sla-profile-1"
            host-limits
                overall 10
```

If the configured host-limit is reached for a subscriber, access is denied for a new host, and an event is generated.

```
80 2015/07/02 09:44:28.30 UTC WARNING: PPPOE #2001 vprn1 PPPoE session failure
"PPPoE session failure on SAP 1/1/1:1 in service 1 -
 [00:00:67:14:01:02,12,user1@domain1] sla-profile sla-profile-1 : host-
limit overall (1) exceeded for subscriber PPPoE-host-user1@domain1 on SAP 1/1/1:1 "
```

Note: An optional command **remove-oldest** can be specified. In this case the new host is accepted and the old one will be removed.

```
configure
    subscriber-mgmt
        sla-profile "sla-profile-1"
                host-limits
                    remove-oldest
```

multi-sub-sap

This parameter defines the maximum number of subscribers (dynamic and static) that can be simultaneously active on this SAP.

When omitted, a single PPPoE session per sap is allowed (no multi-sub-sap).

```
configure
    service
        vprn 1
            subscriber-interface "sub-int-1"
                group-interface "group-int-1"
                    sap 1/1/1:1
                        sub-sla-mgmt
                            multi-sub-sap 100
```

A trap is generated when trying to instantiate a new PPPoE session while the configured number of the multi-sub-sap is reached.

```
123 2015/07/02 09:57:40.80 UTC WARNING: PPPOE #2001 vprn1 PPPoE session failure
"PPPoE session failure on SAP 1/1/1:1 in service 1 -
 [00:00:67:14:01:03,1,user3@domain1] Number of subscribers exceeds the configured mu
lti-sub-sap limit (2)"
```

Redundancy

Redundancy for PPPoE sessions can be used for load balancing the sessions between the 2 BSRs. PADO-delays (which can come from RADIUS, LUDB, and policy) are used to achieve that.

The redundant BSRs need different IP subnets, and upon failure the PPP sessions will need to be re-established.

Because PADI messages are broadcast on a multi-access network, all BSRs on that network will reply with a PADO to the initiator.

The PADR and PADS are sent in unicast to the MAC address from the first received PADO message.

In order to allow control over the NAS/BSR selection for a given PPPoE session the 7750 and 7710 offer the ability to delay the PADO message. Due to the fact that PPPoE clients select the NAS/BSR for further communication based on the first PADO message that arrives, this functionality provides control over the NAS/BSR who gets selected for a given PPPoE session.

On top if for some reason a NAS/BSR, without PADO delay configured in the PPPoE policy, does not reply on PADI messages to the client another NAS/BSR with a PADO delay configured will reply based on the time configured and ultimately the PPPoE session will be established with the PADO delayed NAS/BSR.

Figure 135 **Pado-Delay Scenario**

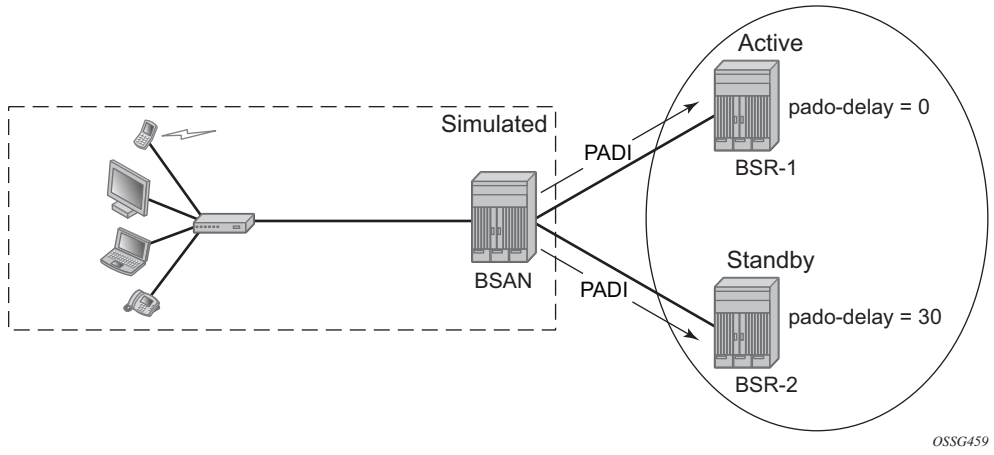

To check the PADO delay value.

```
*A:BSR-1# show subscriber-mgmt ppp-policy "ppp-policy-1"
===============================================================================
PPP Policy "ppp-policy-1"
===============================================================================
Description          : (Not Specified)
Last Mgmt Change     : 07/01/2015 10:03:28
PPP-mtu              : N/A                      Force PPP-mtu >1492  : No
Keepalive Interval   : 30s                      Keepalive Multiplier : 3
Disable AC-Cookies   : Yes                      PADO Delay           : 3000msec
Max Sessions-Per-Mac : 63                       Reply-On-PADT        : Yes
Allow Same CID       : No                       Re-establish Session : Disabled
PPP-Authentication   : pref-CHAP                PPP-CHAP Challenge   : 32 - 64
PPP-Init-Delay (ms)  : 0                        IPCP negotiate subnet: No
Unique SIDs-Per-SAP  : disabled                 Reject-Disabled-Ncp  : No
Ignore-Magic-Num     : No                       Session Timeout      : unlimited
PADO AC-Name         : (Not Specified)
Default username     : (Not Specified)
Default password     : (Not Specified)
```

```
--- snipped ---
*A:BSR-1#
```

Another option to achieve redundancy is through Multi Chassis Synchronization (MCS), but that is beyond the scope of this chapter.

High Availability

The PPPoE session state is HA: the session state is synchronized to the standby CPM. When the active CPM fails, all PPPoE hosts stay active without service interruption.

Conclusion

This chapter provides configuration and troubleshooting commands for PPPoE hosts in a Layer 3 Routed CO (IES/VPRN subscriber interface) context.

ESMv6: IPoE Dual Stack Hosts

This chapter describes IPoE dual stack hosts for ESMv6 configurations.

Topics in this chapter include:

Applicability

This chapter describes ESMv6: IPoE dual stack hosts and is based on SR OS 8.0R4.

This chapter focuses on IPoE IPv6. IPv4 configuration is shown for completeness and is described in more detail in IPv4 DHCP Hosts.

Pre-requisites

Configuring IPoE dual stack hosts for ESMv6 are dependent on the following.

- IOM3-XP or IMM required for subscriber and network interfaces
- Chassis-mode C or higher
- Routed CO (IES/VPRN service) with Enhanced Subscriber Management (ESM)
- Routed Gateway (RG) in the home

Summary

In this chapter, the configuration, operation and troubleshooting of IPoE dual stack hosts in a routed home gateway environment is explained. Focus is on the Enhanced Subscriber Management for IPv6 (ESMv6) part where DHCPv6 is used for IPv6 address assignment. In the BNG, authentication, authorization and IPv6 prefix configuration for an IPoE IPv6 host can be done by a local user database or RADIUS.

Overview

IPoE Dual Stack Hosts

An IPoE dual stack subscriber may support both IPv4 and IPv6 simultaneously. The dual stack hosts share a common subscriber identification policy and have a common SLA- and Subscriber-profile.

IPoE IPv4 and IPv6 hosts operate independently as they are set up through different protocols, DHCPv4 and DHCPv6 respectively. Table 18 and Table 19 show the valid combinations of authentication, authorization and address assignment in the BNG for both address families.

Table 18 **Valid Combinations for RADIUS Authenticated Hosts**

	Authentication and authorization (Subscriber ID and strings)	**IP address assignment (prefix, prefix length, gateway, DNS, etc.)**
IPv6 host	RADIUS	RADIUS
IPv4 host	Static host	Static host
	RADIUS	RADIUS or DHCPv4

Table 19 **Valid Combinations for LUDB Authenticated Hosts**

	Authentication and authorization (Subscriber ID and strings)	**IP address assignment (prefix, prefix length, gateway, DNS)**
IPv6 host	LUDB	LUDB
IPv4 host	Static host	Static host
	Python/DHCPv4	DHCPv4
	SAP defaults	DHCPv4
	LUDB	LUDB or local DHCPv4 server

For an IPoE dual stack subscriber, up to three different types of subscriber hosts can be instantiated.

Figure 136 **IPoE Dual Stack Subscriber Hosts**

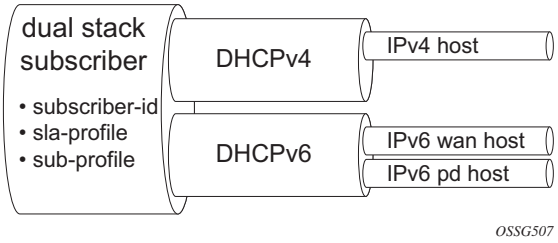

IPoE dual stack subscriber hosts are initially supported in a vlan/subscriber (1:1) routed CO model and with a Routed Gateway (RG). The IPv6 IPoE hosts must support DHCPv6.

Dual Stack IPoE Routed Gateway Service

In the dual stack IPoE Routed Gateway service, the RG in the home network obtains an IPv4 address through the DHCPv4 protocol and an IPv6 Prefix Delegation (PD) prefix and/or wan-host IPv6 address through the DHCPv6 protocol. The Broadband Network Gateway (BNG) authenticates and authorizes both sessions independently.

In the home network, the dual stack RG performs Network Address Translation (NAT) for IPv4, using the assigned IPv4 address as outside address. A global unique IPv6 prefix per subscriber is delegated by the BNG to the RG for use in the home network. The RG can use Stateless Address Auto Configuration (SLAAC) or DHCPv6 to allocate IPv6 addresses from this so called Prefix Delegation (PD) prefix to the devices in the home network. The wan-host IPv6 address is used by the RG on the wan side (network facing). In case of an unnumbered RG, no wan-host address is obtained.

Figure 137 **Dual Stack IPoE Routed Gateway Service**

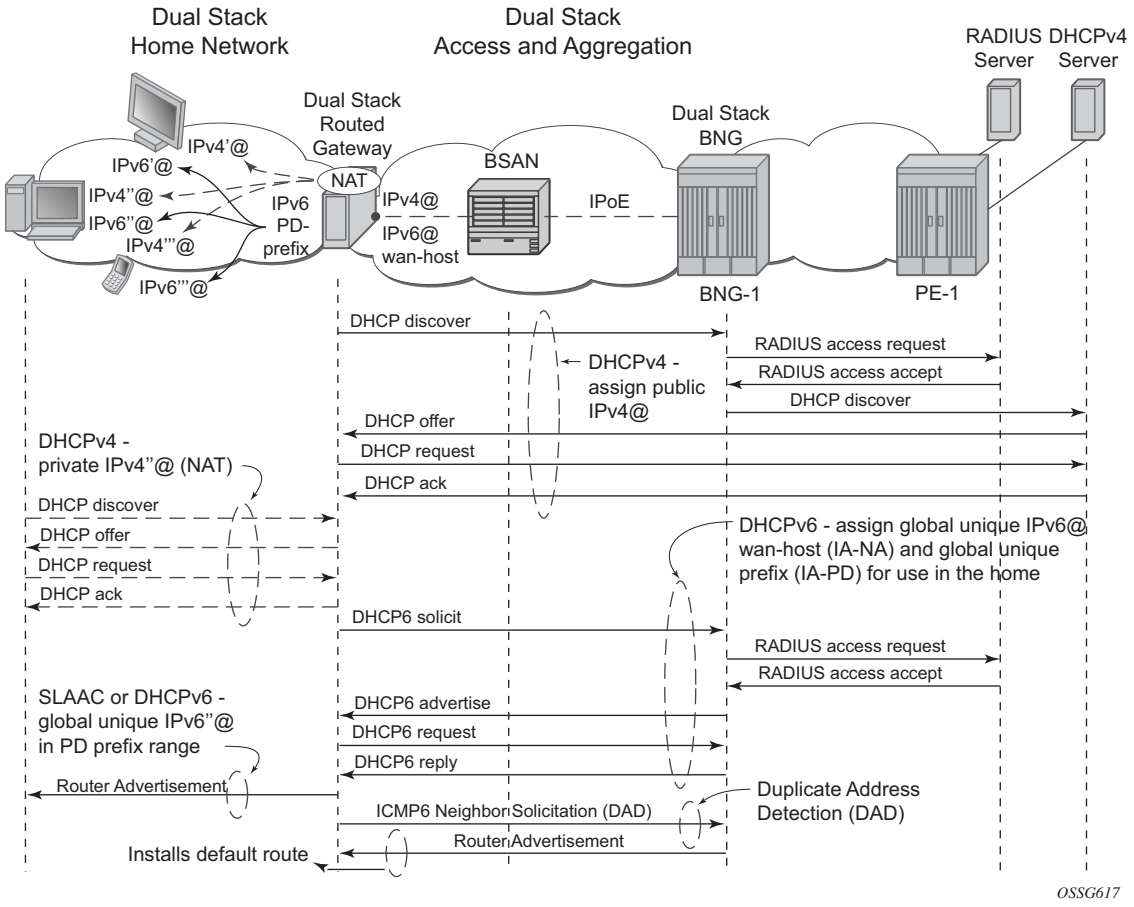

Recap of the DHCPv6 Protocol

The Dynamic Host Configuration Protocol for IPv6 (DHCPv6) is defined in RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*. The protocol enables DHCPv6 servers to pass configuration parameters such as IPv6 network addresses to IPv6 nodes.

DHCPv6 uses the Identity Association (IA) option to assign IPv6 addresses or prefixes. Two different IA types will be used in this section:

- Identity Association for Non-temporary Address (IA-NA) defined in RFC 3315. Used for wan-host IPv6 address assignment.

```
Option : IA_NA (3), Length : 40
  IAID : 1
  Time1: 1800 seconds
```

```
        Time2: 2880 seconds
        Option : IAADDR (5), Length : 24
          Address : 2001:DB8:B001:101::1
          Preferred Lifetime : 3600 seconds
          Valid Lifetime     : 86400 seconds
```

- Identity Association for Prefix Delegation (IA-PD), defined in RFC3633. Used for prefix delegation assignment (for an explanation on prefix delegation, see Prefix Delegation)

```
Option : IA_PD (25), Length : 41
       IAID : 1
       Time1: 1800 seconds
       Time2: 2880 seconds
       Option : IAPREFIX (26), Length : 25
         Prefix : 2001:DB8:A001:100::/56
         Preferred Lifetime : 3600 seconds
         Valid Lifetime     : 86400 seconds
```

The DHCPv6 lease process is outlined in Figure 138 and Figure 139.

Figure 138 **DHCPv6 Lease Process (Part A)**

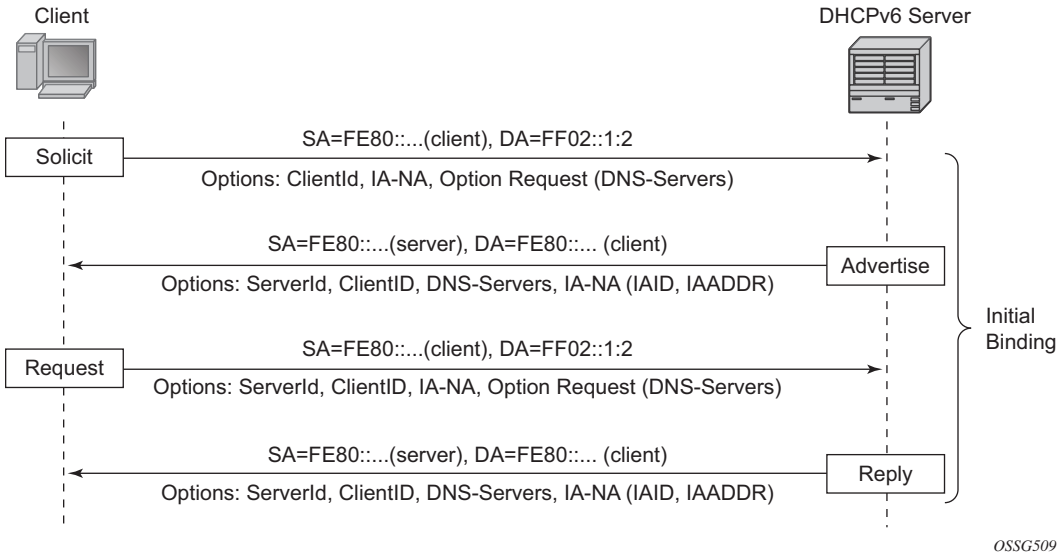

A DHCPv6 client, sends a SOLICIT message to locate servers to the All DHCPv6 Relay Agents and Servers link-scoped multicast address (FF02::1:2), using its link-local address as source address. The DHCPv6 client includes in the SOLICIT message its ClientID, Identity Associations (IA) to request IPv6 address or prefix allocation and optionally an Option Request option.

Any on-link DHCPv6 server responds with a unicasted ADVERTISE message using the link local addresses. The server includes in the ADVERTISE message the ClientID, its ServerID, IPv6 addresses and/or prefixes in Identity Associations (IA) and options containing the requested configuration parameters.

The DHCPv6 client selects an ADVERTISE message and sends a REQUEST message to the All DHCPv6 Relay Agents and Servers link-scoped multicast address. It includes its ClientID, the ServerID of the corresponding DHCPv6 server, Identity Associations (IA) to request IPv6 address or prefix allocation and optionally an Option Request option.

Upon receipt of a valid REQUEST message, the DHCPv6 server with corresponding ServerID, sends a unicast REPLY message using the link local addresses. The REPLY contains the ClientID and ServerID, IPv6 addresses and/or prefixes in Identity Associations (IA) and options containing the requested configuration options.

The DHCPv6 client should perform Duplicate Address Detection (DAD) on the addresses in any IA it received in the REPLY before using that address for traffic.

Figure 139 **DHCPv6 Lease Process (Part B)**

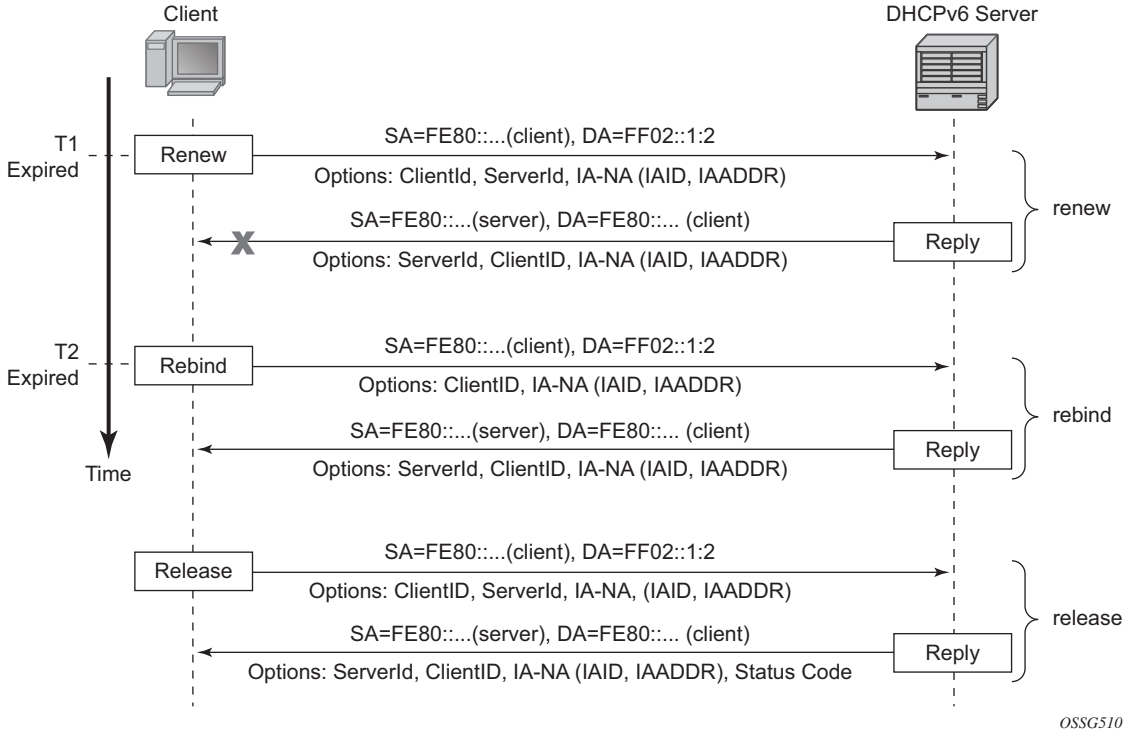

Upon expiration of the renew timer T1 associated with the Identity Association option, the DHCPv6 client sends a RENEW to the All DHCPv6 Relay Agents and Servers link-scoped multicast address to request an extension of the lifetime of an address. It includes its ClientID, the ServerID of the DHCPv6 server that originally provided the address and Identity Associations (IA) containing the IPv6 address or prefix for which an extension of the lifetime is requested.

Upon expiration of the rebind timer T2 associated with the Identity Association option (no response received to the RENEW), the DHCPv6 client sends a REBIND to the All DHCPv6 Relay Agents and Servers link-scoped multicast address to request an extension of the lifetime of an address. It includes its ClientID and Identity Associations (IA) containing the IPv6 address or prefix for which n extension of the lifetime is requested.

If a DHCPv6 client no longer uses one or more of the assigned addresses or prefixes, it sends a RELEASE message to the server that assigned the address or prefix. The server acknowledges with a REPLY message and includes a status code (for example, success).

If the DHCPv6 server sends a Server Unicast Option, then the DHCPv6 client should unicast the REQUEST, RENEW, RELEASE and DECLINE messages to the server using the IPv6 address specified in the option. The 7750 SR DHCPv6 proxy-server does not include the Server Unicast Option.

The DHCPv6 client should perform Duplicate Address Detection (DAD) on each of the addresses assigned through DHCPv6, before using that address for traffic. The DHCPv6 client uses Neighbor Solicitation for this purpose as described in RFC 4862, *IPv6 Stateless Address AutoConfiguration*.

Unlike DHCPv4, DHCPv6 does not provide a default route. In IPv6, default routers are learned via Router Advertisements (see Enable Router Advertisements).

Prefix Delegation

Prefix Delegation (PD) is a mechanism for automated delegation of IPv6 prefixes using DHCPv6. A delegating router delegates a long-lived IPv6 prefix to a requesting router. The delegating router does not require knowledge about the topology of the links in the network to which the prefixes will be assigned.

Figure 140 **Prefix Delegation**

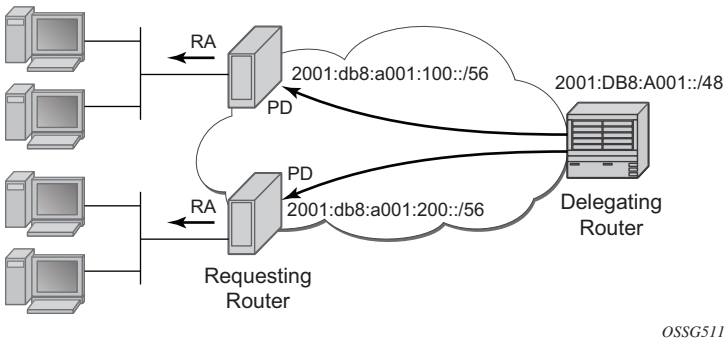

In the context of ESM IPv6, the BNG is acting as delegating router (DHCPv6 server) and the Routed Gateway in the home as requesting router (DHCPv6 client). The DHCPv6 option Identity Association for Prefix Delegation (IA-PD) (Figure 140) is used to assign the IPv6 prefix.

Note that the mechanism through which a requesting router (routed gateway) assigns IPv6 addresses on its interfaces (home network) is arbitrary and can be based upon SLAAC (as shown in Figure 140) or DHCPv6.

Configuration

ESMv6 for IPoE is applicable in a Routed CO environment. The two scenarios below show a minimal configuration to enable dual stack subscribers in a VPRN service context.

Note: ESM IPv6 specific parts are highlighted.

Note: There are no subscriber QoS policies defined (out of scope for this section)

Scenario 1

RADIUS authentication and authorization (later referenced as RADIUS).

```
A:BNG-1# configure subscriber-mgmt
A:BNG-1>config>subscr-mgmt# info
----------------------------------------------
        authentication-policy "radius-1" create
            description "Radius authentication policy"
            password <encrypted password>
            radius-authentication-server
                router "Base"
                server 1 address 172.16.1.1 secret <encrypted secret>
            exit
        exit
        sla-profile "sla-profile-1" create
        exit
        sub-profile "sub-profile-1" create
        exit
        sub-ident-policy "sub-ident-1" create
            sub-profile-map
                use-direct-map-as-default
            exit
            sla-profile-map
                use-direct-map-as-default
            exit
        exit
----------------------------------------------
A:BNG-1>config>subscr-mgmt# exit all


A:BNG-1# configure service vprn 1
A:BNG-1>config>service>vprn# info
----------------------------------------------
            vrf-import "import-1"
            route-distinguisher 64496:1
            auto-bind ldp
            vrf-target export target:64496:1
            subscriber-interface "sub-int-1" create
                address 10.1.255.254/16
                dhcp
                    gi-address 10.1.255.254
                exit
                group-interface "group-int-1" create
                    description "radius authentication and authorization"
                    ipv6
                        router-advertisements
                            managed-configuration
                            no shutdown
                        exit
                        dhcp6
                            proxy-server
                                no shutdown
                            exit
                        exit
                    exit
                    dhcp
                        server 172.16.0.1
                        trusted
                        lease-populate 10
                        no shutdown
                    exit
```

```
                    authentication-policy "radius-1"
                    sap 1/1/2:1 create
                        sub-sla-mgmt
                            sub-ident-policy "sub-ident-1"
                            multi-sub-sap 10
                            no shutdown
                        exit
                    exit
                exit
                ipv6
                    delegated-prefix-len 56
                    subscriber-prefixes
                        prefix 2001:DB8:A001::/48 pd
                        prefix 2001:DB8:B001:100::/56 wan-host
                    exit
                exit
            exit
            service-name "dual-stack"
            no shutdown
-----------------------------------------------
A:BNG-1>config>service>vprn#
```

Scenario 2

Local User Database for authentication and authorization (later referenced as LUDB).

```
*A:BNG-1# configure subscriber-mgmt
*A:BNG-1>config>subscr-mgmt# info
----------------------------------------------
        sla-profile "sla-profile-1" create
        exit
        sub-profile "sub-profile-1" create
            radius-accounting-policy "aaa-policy"
        exit
        sub-ident-policy "sub-ident-1" create
            sub-profile-map
                use-direct-map-as-default
            exit
            sla-profile-map
                use-direct-map-as-default
            exit
            strings-from-option 254
        exit
        local-user-db "ludb-1" create
            dhcp
                match-list mac
                host "host-1" create
                    host-identification
                        mac 00:0a:bc:00:00:01
                    exit
                    address gi-address
                    identification-strings 254 create
                        subscriber-id "sub-1"
                        sla-profile-string "sla-profile-1"
```

```
                        sub-profile-string "sub-profile-1"
                    exit
                    options
                        subnet-mask 255.255.0.0
                        default-router 10.1.255.254
                    exit
                    ipv6-address 2001:DB8:B001:101::1
                    ipv6-prefix 2001:DB8:A001:100::/56
                    no shutdown
                exit
            exit
            no shutdown
        exit
-------------------------------------------
*A:BNG-1>config>subscr-mgmt# exit all


A:BNG-1# configure service vprn 1
A:BNG-1>config>service>vprn# info
-------------------------------------------
            dhcp
                local-dhcp-server "dhcp-s1" create
                    user-db "ludb-1"
                    use-gi-address
                    pool "pool-1" create
                        subnet 10.1.0.0/16 create
                            options
                                subnet-mask 255.255.0.0
                                default-router 10.1.255.254
                            exit
                            address-range 10.1.0.1 10.1.0.255
                        exit
                    exit
                    no shutdown
                exit
            exit
            vrf-import "import-1"
            route-distinguisher 64496:1
            auto-bind ldp
            vrf-target export target:64496:1
            interface "dhcp-s1"  create
                address 192.0.2.1/32
                local-dhcp-server "dhcp-s1"
                loopback
            exit
            subscriber-interface "sub-int-1" create
                address 10.1.255.254/16
                dhcp
                    gi-address 10.1.255.254
                exit
                group-interface "group-int-2" create
                    description "Local user database authentication and authorizatio
n"
                    ipv6
                        router-advertisements
                            managed-configuration
                            no shutdown
                        exit
                        dhcp6
```

```
                            user-db "ludb-1"
                            proxy-server
                                no shutdown
                            exit
                        exit
                    exit
                    dhcp
                        server 192.0.2.1
                        trusted
                        lease-populate 10
                        no shutdown
                    exit
                    sap 1/1/2:2 create
                        sub-sla-mgmt
                            sub-ident-policy "sub-ident-1"
                            multi-sub-sap 10
                            no shutdown
                        exit
                    exit
                exit
                ipv6
                    delegated-prefix-len 56
                    subscriber-prefixes
                        prefix 2001:DB8:A001::/48 pd
                        prefix 2001:DB8:B001:100::/56 wan-host
                    exit
                exit
            exit
            service-name "dual-stack"
            no shutdown
-----------------------------------------------
A:BNG-1>config>service>vprn#
```

Configuring IPv6 Subscriber Prefixes

Applies to both scenarios RADIUS and LUDB.

IPv6 subscriber prefixes must be defined at the **subscriber-interface** <*sub-int-name*> **ipv6 subscriber-prefixes** context. Three types of prefixes can be configured:

- **wan-host** — Prefix from which the IPv6 addresses are assigned that are to be used on the Routed Gateway WAN interface (network facing).
- **pd** — Prefix from which the IPv6 Prefix Delegation prefixes are assigned that are to be used by the Routed Gateway for allocation in the home network (LAN interfaces).
- **pd wan-host** (both) — Prefix from which both IPv6 addresses (wan-host) and IPv6 Prefix Delegation prefixes (pd) can be assigned. This requires that the delegated prefix length is set to 64 bits.

A subscriber prefix length must be between /32 and /63.

Subscriber prefixes are subnetted in fixed length subnets that are assigned to subscriber hosts:

- /64 for **wan-host** subscriber prefixes

  A /128 IPv6 address is assigned to the subscriber host. Broadband Forum standards requires a /64 prefix per subscriber even when used for WAN interfaces and thus the full
  /64 subnet gets associated with the subscriber host [ref. WT-177 - IPv6 in the context of TR-101]. Two subscriber hosts cannot get an IPv6 address from the same /64 subnet.
- /delegated-prefix-len (/48..64) for **pd** subscriber prefixes

  The delegated prefix length is configured in the **subscriber-interface** <*sub-int-name*> **ipv6** context. The recommended value by Broadband Forum standards is /56 (default =
  /64) [ref. WT-177 - IPv6 in the context of TR-101]. The configured length applies to all **pd** subscriber prefixes on a subscriber-interface.

Table 20 provides an overview of the subscriber-prefix parameters that apply:

Table 20 **Applicable Subscriber-Prefix Parameters**

Subscriber prefix type	Subscriber prefix length	DHCPv6 option	Must be subnetted as
wan-host	/32..63	IA-NA	/64 (assigned as /128)
pd	/32..63 (*)	IA-PD	/delegated-prefix-len

(*) must be smaller than configured delegated prefix length

Enable DHCPv6 Proxy Server

Applies to RADIUS and LUDB scenarios.

An IPv6 IPoE subscriber host initiates a DHCPv6 session to request its configuration data (IPv6 addresses and/or IPv6 PD prefixes, DNS servers). Upon receipt of a DHCPv6 SOLICIT message, the BNG authenticates the IPv6 subscriber host and obtains its configuration information from a RADIUS server or local user database. A DHCPv6 proxy server in the BNG maintains the DHCPv6 session with the IPv6 IPoE subscriber host.

The DHCPv6 proxy server must be enabled in the **subscriber-interface** <*sub-int-name*> **group-interface** <*group-int-name*> **ipv6 dhcp6 proxy-server** context. The default is **shutdown**.

```
service
    vprn 1 customer 1 create
        subscriber-interface "sub-int-1" create
            group-interface "group-int-1" create
                ipv6
                    dhcp6
                        proxy-server
                            renew-timer 1800           # default
                            rebind-timer 2880          # default
                            valid-lifetime 86400       # default
                            preferred-lifetime 3600    # default
                            client-applications dhcp   # default
                            no shutdown
                        exit
```

When enabled, the DHCPv6 proxy server by default allows IPv6 IPoE hosts to authenticate (configured with client-applications dhcp. Additionally, you can enable support for IPv6 PPPoE hosts. See ESMv6: PPPoE Dual Stack Hosts.

A number of timers associated with IPv6 addresses and IPv6 prefixes within DHCPv6 Identity Associations can be configured in the DHCPv6 proxy server.

RFC 4862 defines two timers associated with graceful degradation of address bindings:

- Preferred lifetime — The length of time that a valid address is preferred (the time until deprecation). When the preferred lifetime expires, the address becomes deprecated and its use should be discouraged for new sessions.
- Valid lifetime — The length of time an address remains in the valid state (the time until invalidation). The valid lifetime must be greater than or equal to the preferred lifetime. When the valid lifetime expires, the address becomes invalid.

RFC 3315 (DHCPv6) defines two timers associated with an Identity Association (IA) option that give the servers explicit control over when a client recontacts the server about a specific IA:

- T1 (renew) — The time at which the client contacts the server from which the addresses/prefix in the IA were obtained to extend the lifetimes of the addresses/prefix assigned to the IA
- T2 (rebind) — The time at which the client contacts any available server to extend the lifetimes of the addresses/prefixes assigned to the IA;

These timers are common for all DHCPv6 sessions in a group-interface and cannot be configured from RADIUS nor local user database.

Figure 141 **IPv6 Address/Prefix Timers**

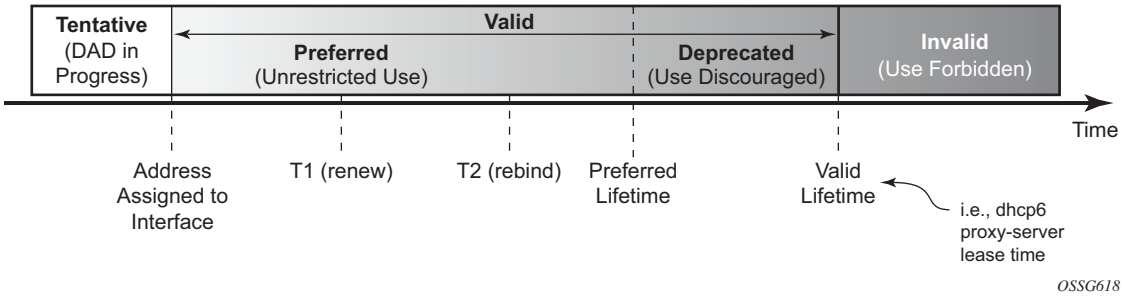

When violating the following rule, the default timers will be used:

Table 21 **Timer Parameters**

Timer	Use	Default	Range
T1	Renew timer	1800s (30 min)	0..604800s (7 days)
T2	Rebind timer	2880s (48 min)	0..1209600s (14 days)
preferred-lifetime		3600s (1hr)	300..4294967295s
valid-lifetime	DHCPv6 lease time	86400s (24 hrs)	300..4294967295s

If the DHCPv6 lease is not renewed by the client before the DHCPv6 lease timer expires, then the subscriber host is deleted from the system. In other words, beyond the valid lifetime, subscriber traffic from/to the associated IPv6 addresses is dropped.

Enable Router Advertisements

Applies to both scenarios RADIUS and LUDB.

In IPv6, default routers are automatically installed via the router discovery mechanism. Unsolicited Router Advertisements (RA) must explicitly be enabled on a group interface. The default is **shutdown**.

```
service
    vprn 1 customer 1 create
        subscriber-interface "sub-int-1" create
            group-interface "group-int-1" create
                ipv6
```

```
router-advertisements
    managed-configuration
    no shutdown
exit
```

Note that the managed-configuration flag is set for consistency only. It tells the hosts that addresses are available by DHCPv6. However, as described in the Security section later (see Security), the host cannot rely on this flag because DHCPv6 must be initiated by the host before the BNG sends RAs.

Additional parameters that can be configured with respect to the router advertisements (defaults are shown):

```
service
    vprn 1 customer 1 create
        subscriber-interface "sub-int-1" create
            group-interface "group-int-1" create
                ipv6
                    router-advertisements
                        shutdown
                        current-hop-limit 64
                        no managed-configuration
                        max-advertisement 1800
                        min-advertisement 900
                        no mtu
                        no other-stateful-configuration
                        prefix-options
                            no autonomous
                            preferred-lifetime 3600
                            valid-lifetime 86400
                        exit
                        reachable-time 0
                        retransmit-time 0
                        router-lifetime 4500
                    exit
```

Table 22 **Router Advertisements Parameters**

Parameter	**Description (RFC 4861,** *Neighbor Discovery for IP version 6 (IPv6)***)**	**Value Range (default)**
current-hop-limit	The default value that should be placed in the Hop Count field of the IP header for outgoing IP packets. A value of zero means unspecified (by this router); the RG picks its own value.	0..255 (64)
managed-configuration	Managed address configuration flag. When set, it indicates that addresses are available through DHCPv6	(no)

Table 22 **Router Advertisements Parameters (Continued)**

Parameter	Description (RFC 4861, *Neighbor Discovery for IP version 6 (IPv6)*)	Value Range (default)
max-advertisement	Unsolicited Router Advertisements are not strictly periodic: the interval between subsequent transmissions is randomized to reduce the probability of synchronization with the advertisements from other routers on the same link. Whenever a multicast advertisement is sent from an interface, the timer is reset to a uniformly distributed random value between the interface's configured MinRtrAdvInterval and MaxRtrAdvInterval.	900..1800 s (1800)
min-advertisement		900..1350 s (900)
mtu	Routers can advertise an MTU for hosts to use on the link.	1280..9212 bytes (no)
other-stateful-configuration (not applicable for IPoE)	Other configuration flag. When set, it indicates that other configuration information is available through DHCPv6. (DNS). Can be ignored if managed address configuration flag is enabled	(no)
prefix-options: autonomous (not applicable for IPoE)	Autonomous address-configuration flag. When set indicates that this prefix can be used for stateless address autoconfiguration (SLAAC)	(no)
prefix-options: preferred-lifetime (not applicable for IPoE)	The length of time in seconds that addresses generated from the prefix via stateless address autoconfiguration (SLAAC) remain preferred	0..4294967295 (3600)
prefix-options: valid-lifetime (not applicable for IPoE)	The length of time in seconds that the prefix is valid for the purpose of on-link determination. (also used by SLAAC)	0..4294967295 (86400)
reachable-time	The time that a node assumes a neighbor is reachable after having received a reachability confirmation. Used by the Neighbor Unreachability Detection algorithm. A value of zero means unspecified (by this router); the RG picks its own value.	0..3600000 ms (0)
retransmit-time	The time between retransmitted Neighbor Solicitation messages. Used by address resolution and the Neighbor Unreachability Detection algorithm. A value of zero means unspecified (by this router); the RG picks its own value.	0..1800000 ms (0)
router-lifetime	The lifetime associated with the default router in units of seconds.	2700..9000 s (4500)

RADIUS Authentication and Authorization

Applies to scenario 1 RADIUS only.

The RADIUS authentication and authorization configuration for IPoE IPv6 subscriber host is no different from an IPv4 subscriber host:

```
subscriber-mgmt
    authentication-policy "radius-1" create
        description "Radius authentication policy"
        password <hashed password> hash2
        radius-authentication-server
            router "Base"
            server 1 address 172.16.1.1 secret <hashed secret> hash2
        exit
    exit

    vprn 1 customer 1 create
        subscriber-interface "sub-int-1" create
            group-interface "group-int-1" create
                authentication-policy "radius-1"
```

Additional RADIUS AVPs that are applicable for IPoE IPv6 subscriber hosts are listed in Table 23.

Table 23 **RADIUS AVPs**

RADIUS AVP	Type	Purpose
Alc-IPv6-Address [26-6527-99]	ipv6addr	maps to IA_NA of DHCPv6 (RG WAN interface address)
Alc-Ipv6-Primary-Dns [26-6527-105]	ipv6addr	maps to DNS Recursive Name Server option (RFC 3646, *DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*) in DHCPv6
Alc-Ipv6-Secondary-Dns [26-6527-106]	ipv6addr	maps to DNS Recursive Name Server option (RFC 3646) in DHCPv6
Delegated-IPv6-Prefix [123]	ipv6prefix	maps to IA_PD for prefix delegation (RFC 3633, *IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6*) in DHCPv6

A sample FreeRADIUS users record to authenticate a dual stack IPoE subscriber:

```
00:0a:bc:00:00:01   Auth-Type := Local, Cleartext-Password := "password"
                    Alc-Subsc-ID-Str = "sub-1",
                    Alc-Subsc-Prof-Str = "sub-profile-1",
                    Alc-SLA-Prof-Str = "sla-profile-1",
                    Alc-IPv6-Address = 2001:db8:b001:101::1,
                    Delegated-IPv6-Prefix = 2001:db8:a001:100::/56,
                    Alc-Ipv6-Primary-DNS = 2001:db8:dddd:1::1,
                    Alc-Ipv6-Secondary-DNS = 2001:db8:dddd:2::1
```

Note the FreeRADIUS Server 2.0.0 and greater has full support for both IPv6 attributes and IPv6 network packets.

The IPv6 address/prefix related timers can be configured in the **dhcp6 proxy-server** context (see Enable DHCPv6 Proxy Server).

Local User Database Authentication and Authorization

Applies to scenario 2 LUDB only.

The configuration example below focuses on the IPv6 host configuration. The details for local user database host matching and IPv4 host specific parameters are out of scope for this section.

```
subscriber-mgmt
    local-user-db "ludb-1" create
        dhcp
            match-list mac
            host "host-1" create
                host-identification
                    mac 00:0a:bc:00:00:01
                exit
                address gi-address                      # IPv4 host
                identification-strings 254 create
                    subscriber-id "sub-1"
                    sla-profile-string "sla-profile-1"
                    sub-profile-string "sub-profile-1"
                exit
                options
                    subnet-mask 255.255.0.0             # IPv4 host
                    default-router 10.1.255.254         # IPv4 host
                exit
                ipv6-address 2001:DB8:B001:101::1       # IPv6 host
                ipv6-prefix 2001:DB8:A001:100::/56      # IPv6 host
                no shutdown
            exit
        exit
        no shutdown
    exit
exit

    vprn 1 customer 1 create
        subscriber-interface "sub-int-1" create
            group-interface "group-int-2" create
                ipv6
                    dhcp6
                        user-db "ludb-1"
```

Next to the identification strings that are common between IPv4 and IPv6 hosts, there are two specific IPv6 host related parameters to be configured:

Table 24 **Local User Database Parameters**

local-user-db CLI parameter	Purpose
ipv6-address	Maps to IA_NA of DHCPv6 (RG WAN interface address)
ipv6-prefix	Maps to IA_PD for prefix delegation (RFC 3633) in DHCPv6

The IPv6 address/prefix related timers can be configured in the **dhcp6 proxy-server** context (see Enable DHCPv6 Proxy Server).

Note that DNSv6 server information cannot be configured in the local user database scenario. The DNSv6 server information should either be manually configured on the host or a DNSv4 server should be used instead.

DHCP and DHCP6 Lease State

Applies to both scenarios RADIUS and LUDB.

The DHCP lease state is an internal database structure that keeps track of the DHCP host states. The DHCP lease state enables subscriber management functions (for example, per subscriber QoS and accounting) and security functions (for example, dynamic anti-spoof filtering) on the DHCP host.

The DHCP lease information for a specific host is extracted from the DHCPv4 ack message in case of DHCPv4 and from the DHCPv6 reply message in case of DHCPv6

Typical information stored in the DHCP lease state includes (partial table; additional data can be stored for managed SAPs, wholesale-retail).

Table 25 **DHCP Lease State Information**

Parameter	Comment
Service ID	Service where the DHCP host is connected.
IP Address	IPv4 or IPv6 address of the DHCP host.
Client HW Address	Ethenet MAC address of the DHCP host.

Table 25 **DHCP Lease State Information (Continued)**

Parameter	Comment
Subscriber-interface (Routed CO only)	Subscriber interface name where the DHCP host is instantiated.
Group-interface (Routed CO only)	Group interface name where the DHCP host is instantiated.
SAP	SAP where the DHCP hosts is connected.
Remaining Lifetime	The remaining time before the DHCP host is deleted from the system (updated each time a DHCP renew/rebind occurs).
Persistence Key	Lookup key for this host in the persistency file.
Sub-Ident	ESM: Subscriber ID of the DHCP host.
Sub-Profile-String	ESM: Subscriber profile string of the DHCP host.
SLA-Profile-String	ESM: SLA profile string of the DHCP host.
App-Profile-String	ESM: Application profile string of the DHCP host.
Lease ANCP-String	ESM: ANCP string for this DHCP host.
Lease Int Dest Id	ESM: Internal destination ID for this DHCP host.
Category-Map-Name	ESM: Volume and Time based accounting.
Dhcp6 ClientId (DUID)	DHCPv6 client unique identifier.
Dhcp6 IAID	Identity Association ID chosen by the client.
Dhcp6 IAID Type	Identity Association type: prefix (PD) or non-temporary (wan-host).
Dhcp6 Client Ip	Link local IPv6 address of the host.
Sub-Ident origin	ESM: Origin for the Subscriber ID for this host (None, DHCP, RADIUS).
Strings origin	ESM: Origin for the ESM strings for this host (None, DHCP, RADIUS).
Lease Info origin	ESM: Origin for the IP configuration for this host (None, DHCP, RADIUS).
Ip-Netmask	The IP netmask for this DHCP host.
Broadcast-Ip-Addr	The broadcast IP address for this host.
Default-Router	The default gateway for this host.
Primary-Dns	The primary DNS server for this host.
Secondary-Dns	The secondary DNS server for this host.

Table 25 **DHCP Lease State Information (Continued)**

Parameter	Comment
Primary-Nbns	The primary NetBIOS name server for this host.
Secondary-Nbns	The secondary NetBIOS name server for this host.
ServerLeaseStart	Time and date that the lease for this host started (first DHCP ack received).
ServerLastRenew	Time and date that the lease for this host was last renewed.
ServerLeaseEnd	Time and date that the lease for this host will expire.
Session-Timeout	Lease time specified by the DHCP server.
DHCP Server Addr	IP address of the DHCP server that allocated the lease for this host.
Circuit Id	DHCP Relay Agent information option 82 Circuit ID content.
Remote Id	DHCP Relay Agent information option 82 Remote ID content.
RADIUS User-Name	ESM: Username used in the RADIUS authentication access request.

DHCPv4 lease state population is enabled by default on a group-interface with DHCP configured as **no shutdown**. The number of DHCPv4 leases allowed on each SAP of the group-interface must be configured with the **lease-populate** option (by default a single DHCPv4 host is allowed on each SAP of the group-interface).

DCHPv6 lease state population is enabled by default on a group-interface with DHCP6 proxy-server configured as **no shutdown**. The number of DHCPv6 leases (hosts) cannot be limited per group-interface.

```
configure
    service
        vprn 1 customer 1 create
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    ipv6
                        dhcp6
                            proxy-server
                                no shutdown
                            exit
                        exit
                    exit
                    dhcp
                        server 172.16.0.1
                        trusted
                        lease-populate 10
                        no shutdown
                    exit
```

To check the DHCPv4 or DHCPv6 lease state for a particular service, use the following commands (detailed output as well as additional output filtering is available):

```
*A:BNG-1# show service id 1 dhcp lease-state ?
  - 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>]}] [detail]


*A:BNG-1# show service id 1 dhcp6 lease-state detail
===============================================================================
DHCP lease states for service 1
===============================================================================
Service ID           : 1
IP Address           : 2001:DB8:A001:100::/56
Client HW Address    : 00:0a:bc:00:00:01
Subscriber-interface : sub-int-1
Group-interface      : group-int-1
SAP                  : 1/1/2:1
Remaining Lifetime   : 23h59m49s
Persistence Key      : 0x0000004d

Sub-Ident            : "sub-1"
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-1"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : ""
Category-Map-Name    : ""
Dhcp6 ClientId (DUID): 00010001133ebdd2000c29c851ca
Dhcp6 IAID           : 1
Dhcp6 IAID Type      : prefix
Dhcp6 Client Ip      : FE80::20A:BCFF:FE00:1
Primary-Dns          : 2001:DB8:DDDD:1::1
Secondary-Dns        : 2001:DB8:DDDD:2::1

Sub-Ident origin     : Radius
Strings origin       : Radius
Lease Info origin    : Radius

ServerLeaseStart     : 09/02/2010 16:13:11
ServerLastRenew      : 09/02/2010 16:13:11
ServerLeaseEnd       : 09/03/2010 16:13:11
Radius User-Name     : "00:0a:bc:00:00:01"
-------------------------------------------------------------------------------
Service ID           : 1
IP Address           : 2001:DB8:B001:101::1/128
Client HW Address    : 00:0a:bc:00:00:01
Subscriber-interface : sub-int-1
Group-interface      : group-int-1
SAP                  : 1/1/2:1
Remaining Lifetime   : 23h59m49s
Persistence Key      : 0x0000004c

Sub-Ident            : "sub-1"
```

```
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-1"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : ""
Category-Map-Name    : ""
Dhcp6 ClientId (DUID): 00010001133ebdd2000c29c851ca
Dhcp6 IAID           : 1
Dhcp6 IAID Type      : non-temporary
Dhcp6 Client Ip      : FE80::20A:BCFF:FE00:1
Primary-Dns          : 2001:DB8:DDDD:1::1
Secondary-Dns        : 2001:DB8:DDDD:2::1

Sub-Ident origin     : Radius
Strings origin       : Radius
Lease Info origin    : Radius

ServerLeaseStart     : 09/02/2010 16:13:11
ServerLastRenew      : 09/02/2010 16:13:11
ServerLeaseEnd       : 09/03/2010 16:13:11
Radius User-Name     : "00:0a:bc:00:00:01"
-------------------------------------------------------------------------------
Number of lease states : 2
===============================================================================
*A:BNG-1#
```

Operation

An IPoE dual stack subscriber in a numbered Routed Gateway scenario consumes three subscriber host entries:

- IPv4 host — DHCPv4 session based
- IPv6 wan-host — DHCPv6 session based
- IPv6 Prefix Delegation host — DHCPv6 session based

```
*A:BNG-1# show service active-subscribers
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber sub-1 (sub-profile-1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:1/1/2:1 - sla:sla-profile-1
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.1.0.3
                00:0a:bc:00:00:01 N/A       DHCP
2001:DB8:A001:100::/56
                00:0a:bc:00:00:01 N/A       IPoE-DHCP6
```

```
2001:DB8:B001:101::1/128
               00:0a:bc:00:00:01 N/A       IPoE-DHCP6
-------------------------------------------------------------------------------
Number of active subscribers : 1
===============================================================================
*A:BNG-1#
```

The optional **hierarchy** parameter for the active-subscribers display provides a top-down level overview for this subscriber:

```
*A:BNG-1# show service active-subscribers hierarchy
===============================================================================
Active Subscriber hierarchy
===============================================================================
-- sub-1 (sub-profile-1)
   |
   |-- sap:1/1/2:1 - sla:sla-profile-1
   |   |
   |   |-- 10.1.0.3
   |   |   00:0a:bc:00:00:01 - N/A (DHCP)
   |   |
   |   |-- 2001:DB8:A001:100::/56
   |   |   00:0a:bc:00:00:01 - N/A (IPoE-DHCP6)
   |   |
   |   |-- 2001:DB8:B001:101::1/128
   |   |   00:0a:bc:00:00:01 - N/A (IPoE-DHCP6)
   |   |

===============================================================================
A:BNG-1#
```

The total number (sum) of IPv4 and IPv6 hosts per subscriber can be limited in the corresponding sla-profile with the **host-limit** parameter:

```
    subscriber-mgmt
        sla-profile "sla-profile-1" create
            host-limit 3
        exit
```

To display the IPv4/IPv6 routing table for dual stack hosts:

```
A:BNG-1# show router 1 route-table ipv4 protocol sub-mgmt
===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix                                 Type    Proto    Age          Pref
       Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.1.0.3/32                                 Remote  Sub Mgmt 00h01m44s   0
       [group-int-1]                                               0
-------------------------------------------------------------------------------
No. of Routes: 1
===============================================================================
```

```
A:BNG-1#


A:BNG-1# show router 1 route-table ipv6 protocol sub-mgmt
===============================================================================
IPv6 Route Table (Service: 1)
===============================================================================
Dest Prefix                                     Type    Proto    Age          Pref
       Next Hop[Interface Name]                                        Metric
-------------------------------------------------------------------------------
2001:DB8:A001:100::/56                          Remote  Sub Mgmt 00h01m50s   0
       [group-int-1]                                                 0
2001:DB8:B001:101::1/128                        Remote  Sub Mgmt 00h01m50s   0
       [group-int-1]                                                 0
-------------------------------------------------------------------------------
No. of Routes: 2
===============================================================================
A:BNG-1#
```

Troubleshooting

Apart from the show commands in this section, use the following commands to troubleshoot a dual stack host session:

- Default system log:

```
A:BNG-1# show log log-id 99
```

  Use appropriate filtering to reduce the output if needed.

- Debug:

```
debug
    router "1"
        ip
            dhcp                                    # DHCPv4
                detail-level medium
                mode egr-ingr-and-dropped
            exit
            dhcp6                                   # DHCPv6
                mode egr-ingr-and-dropped
                detail-level high   # needed to see the option content
            exit
        exit
        local-dhcp-server dhcp-s1                   # local dhcp server
            detail-level medium
            mode egr-ingr-and-dropped
        exit
    exit
    subscriber-mgmt
        local-user-db ludb-1                        # local user database
            detail all
```

```
        exit
    exit
    radius detail                                 # RADIUS
exit
```

Note that additional filtering (such as only DHCPv6 debug for particular interface) may be needed to prevent flooding of debug messages.

- Protocol statistics:

  DHCPv4 stats:

```
A:BNG-1# show router 1 dhcp statistics
===============================================================================
DHCP Global Statistics (Service: 1)
===============================================================================
Rx Packets                          : 3192
Tx Packets                          : 3177
Rx Malformed Packets                : 0
Rx Untrusted Packets                : 0
Client Packets Discarded            : 0
Client Packets Relayed              : 737
Client Packets Snooped              : 860
Client Packets Proxied (RADIUS)     : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Discarded            : 15
Server Packets Relayed              : 733
Server Packets Snooped              : 847
DHCP RELEASEs Spoofed               : 0
DHCP FORCERENEWs Spoofed            : 0
===============================================================================
A:BNG-1#
```

  DHCPv6 stats:

```
*A:BNG-1# show router 1 dhcp6 statistics
===============================================================================
DHCP6 statistics (Router: 1)
===============================================================================
Msg-type                      Rx             Tx             Dropped
-------------------------------------------------------------------------------
1 SOLICIT                     3              0              0
2 ADVERTISE                   0              3              0
3 REQUEST                     3              0              0
4 CONFIRM                     0              0              0
5 RENEW                       313            0              6
6 REBIND                      0              0              0
7 REPLY                       0              312            0
8 RELEASE                     2              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

-------------------------------------------------------------------------------
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                                    6
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
===============================================================================
A:BNG-1#
```

RADIUS stats:

```
*A:BNG-1# show subscriber-mgmt authentication "radius-1" statistics
===============================================================================
Authentication Policy Statistics
===============================================================================
-------------------------------------------------------------------------------
Policy name                              : radius-1
subscriber packets authenticated         : 16
subscriber packets rejected              : 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      16        0         0         0          0        0
===============================================================================
A:BNG-1#
```

Advanced Topics

Security

Downstream Router Advertisements

When a SAP is bound to a subscriber/group-interface which has IPv6 enabled, there will be no initial downstream Router Advertisement (RA) message sent. If a SAP is shared by multiple subscribers, it would be possible for an unauthenticated host to receive the RA.

Instead the RAs are sent in unicast to allow per-host IPv6 link configuration. This requires the host information (MAC address and link-local IPv6 address) to be known. Hence for IPoE, until a DHCPv6 session is bound, no unsolicited or solicited RAs are send.

Processing of Neighbor Discovery Messages

Processing of Neighbor Discovery messages: Neighbor Advertisements (NA), Neighbor Solicitations (NS) and Router Solicitations (RS).

Neighbor discovery messages are not processed prior to IPoE IPv6 host authentication to avoid DoS attacks consuming CPU resources. This implies that an IPoE host should initiate the DHCPv6 session without link information and knowledge of routers on the link as required by the Broadband Forum standards (ref. TR-124 issue 2 — Functional Requirements for Broadband Residential Gateway Devices). This is not a problem as the DHCPv6 solicit/request messages are sent to a well-known multicast address with direct link-layer mapping.

After DHCP host authentication, Neighbor Discovery messages will not result in a neighbor cache entry. Instead a managed neighbor cache entry is created based on the DHCPv6 lease state. This managed neighbor cache entry cannot be displayed. The above mechanism prevents DoS attacks from poisoning the neighbor cache with bogus entries.

Router advertisements in response to a router solicitation are internally throttled so that they are not sent more often than once every three seconds.

Anti-spoof Filters

For each authenticated IPoE IPv6 host, an anti-spoof filter entry is created that allows upstream traffic with exact match on the tuple {masked source IP, source MAC}. Traffic from unauthenticated hosts is silently dropped.

Managed SAPs

To allow the creation of managed SAPs in a dual stack environment, both DHCPv4 discover and DHCPv6 solicit messages received on a capture SAP should trigger RADIUS authentication:

```
service
    vpls 2 customer 1 create
        sap 1/1/2:* capture-sap create
            trigger-packet dhcp dhcp6
            authentication-policy "radius-1"
        exit
        no shutdown
    exit
```

A full description of the managed SAP functionality is out of the scope of this section.

RADIUS Change of Authorization (CoA)

The only CoA action that is allowed for IPoE IPv6 hosts is a change of ESM strings (SLA-profile, subscriber-profile, application-profile, etc). Creation of a new IPv6 host or forcing a DHCPv6 renew is not supported.

Only a single address attribute (Framed-IP-Address, Delegated-IPv6-Prefix or Alc-IPv6-Address) may be given in a single request. When host-accounting is enabled, only the host specific accounting session IDs (Acct-Session-Id) can be used. This means that to change for example the sla-profile for all three hosts of a dual stack subscriber, three CoA messages should be sent.

A full description of the RADIUS CoA functionality is out of the scope of this section.

Accounting

There are no separate accounting statistics available for IPv4 and IPv6 traffic unless they are mapped in a different Forwarding Class/queue.

In RADIUS accounting, host-accounting could be enabled to see the IPv4 and IPv6 host instantiations separately: an accounting start/stop is generated for each individual subscriber host. The actual accounting data is included in the interim updates and accounting stop message for the sla-profile instance.

A full description of the accounting functionality is out of the scope of this section.

Lease State Persistency

A DHCPv4/DHCPv6 (hereafter referred to as DHCP) session does not have a keep-alive mechanism to detect unavailability. A new DHCP session set-up is only attempted after expiration of the DHCP lease time. A node reboot causing the loss of DHCP lease state and the corresponding anti-spoof filters could therefore result in unacceptable long service outages.

The DHCP lease state can be made persistent across node reboots: DHCP lease state is restored from a persistency file stored on the compact flash file system. As a result, DHCP sessions will only loose connectivity during the time of reboot without being completely disconnected.

To activate the DHCP lease state persistency:

```
configure
    system
        persistence
            subscriber-mgmt
                description "DHCP lease state persistency"
                location cf2:
            exit
        exit
```

A dedicated persistency file will be created on the specified compact flash file system. The file is initialized to store the maximum number of allowed hosts; its size is fixed to avoid file system space problems during operations.

```
*A:BNG-1# file dir cf2:

 Volume in drive cf2 on slot A has no label.
```

```
 Volume in drive cf2 on slot A is formatted as FAT32.

Directory of cf2:\

09/02/2010  01:27p           536871424 submgmt.006
               1 File(s)              536871424 bytes.
               0 Dir(s)              1558183424 bytes free.
```

Each time the DHCP session is renewed, the persistency file is updated together with the lease state. If the file update fails, an event is generated to indicate that persistency can not be guaranteed.

The content of the persistency file may vary between different SR-OS software releases. When upgrading, the persistency file is automatically upgraded to the new format. To downgrade the persistency file to a lower SR-OS release version, use the following command:

```
*A:BNG-1# tools perform subscriber-mgmt downgrade ?
  - downgrade target-version <target> [reboot]

 <target>             : The version you want to downgrade to
                            8.0 (current)  - submgmt.006
                            7.0            - submgmt.005
                            6.0            - submgmt.004
                            5.0            - submgmt.003
                            4.0            - submgmt.pst
 <reboot>             : reboot system after successful conversion
```

The content of the persistency file can be looked at using the following commands:

```
*A:BNG-1# show service id 1 dhcp6 lease-state detail
===============================================================================
DHCP lease states for service 1
===============================================================================
Service ID           : 1
IP Address           : 2001:DB8:A001:100::/56
Client HW Address    : 00:0a:bc:00:00:01
Subscriber-interface : sub-int-1
Group-interface      : group-int-1
SAP                  : 1/1/2:1
Remaining Lifetime   : 23h49m47s
Persistence Key      : 0x0000004d

- - - snip - - -

*A:BNG-1# tools dump persistence submgt record 0x0000004d
-----------------------------------
Persistency File Record
-----------------------------------
Filename    : cf2:\submgmt.006
Key         : 0000004d
Last Update : 2010/09/02 16:13:12 (UTC)
Action      : ADD
Data :
```

```
Host Type      : IpV6 node address
Service ID     : 1
SAP ID         : 1/1/2:1
IP             : 2001:DB8:A001:100::/56
NH MAC         : 00:0a:bc:00:00:01
Created        : 2010/09/02 16:13:11 (UTC)
Session Timeout: 0 (seconds)
Sub-ID         : sub-1
Sub-prof-ID    : sub-profile-1
SLA-prof-ID    : sla-profile-1
App-prof-ID    : NULL
ANCP-Str       : NULL
Int-dest-ID    : NULL
Cat-map-str    : NULL
Sub-Id is def  : NO
Int-dest is def: YES
Address Origin : 1
SubId Origin   : 1
Strings Origin : 1
RADIUS Fallback: NO
Managed routes : None
BgpPrngPlcyAttr: None
Class Attr     : 1 bytes
Radius Username: 00:0a:bc:00:00:01
Pri. IPv6 DNS  : 2001:DB8:DDDD:1::1
Sec. IPv6 DNS  : 2001:DB8:DDDD:2::1
```

Conclusion

This chapter provides configuration, operation and troubleshooting commands for dual stack IPoE subscribers on Routed Gateways. Focus is on the ESMv6 part where DHCPv6 is used for IPv6 address assignment on the RG network interface (wan host) and for allocation of an IPv6 prefix delegation prefix for use in the home network (pd host). In the BNG, authentication, authorization and IPv6 prefix configuration for an IPoE IPv6 host is done by a local user database or RADIUS.

ESMv6: PPPoE Dual Stack Hosts

This chapter describes ESMv6 PPPoE dual stack host configurations.

Topics in this chapter include:

Applicability

This chapter is applicable to SROS routers and was initially based on release 8.0.R.4. The CLI is updated to release14.0.R3.

Pre-requisites:

- IOM3-XP or IMM required for subscriber and network interfaces
- Chassis-mode C or higher
- Routed CO (IES/VPRN service) with Enhanced Subscriber Management (ESM)
- Bridged or routed home gateway

Note: The focus of this chapter is on PPPoE IPv6. IPv4 configuration is shown for completeness.

Overview

PPPoE Dual Stack

A PPPoE dual stack subscriber may support both IPv4 and IPv6 simultaneously. The dual stack hosts share a common subscriber identification policy and have a common sla-profile and subscriber-profile and are linked together through one PPPoE session.

For PPPoE dual stack hosts, one subscriber host is created for IPv4 and another one for the IPv6 address family.

Figure 142 **PPPoE Dual Stack Hosts**

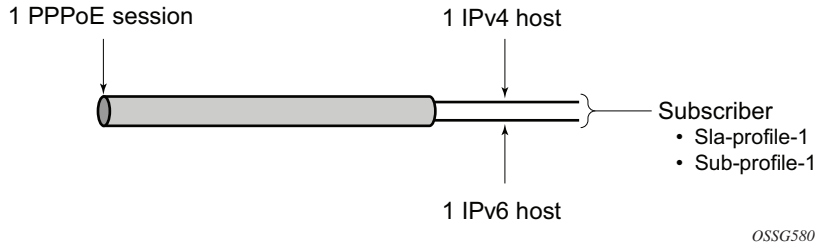

ESM for IPv6 is supported through RADIUS and local user database (LUDB) for authentication, address assignment and authorization.

PPPoE dual stack subscriber-hosts are supported for bridged and routed home gateways.

Dual Stack PPPoE Bridged Gateway Service

In the dual stack PPPoE host service, the PPPoE session is initiated directly from a dual stack device in the home network. PPPoE is used to carry IPv6 and (optionally) IPv4 traffic from the device to the broadband remote access server (BRAS), also called broadband network gateway (BNG).

Unlike the routed gateway application examples (see later), no IPv6 prefix delegation occurs in the bridged gateway service. Instead, a global unicast address prefix (/64) is advertised using Router Advertisements (RAs) directly to the PPPoE interface on the host.

The device addresses are self-assigned through stateless auto configuration (SLAAC). where SLAAC makes use of ICMPv6 router-advertisements to announce these IPv6 prefixes. The SLAAC prefixes have a mandatory length of /64.

This application is targeted at operators who currently use a bridging modem in the customer premises and who want to incrementally add IPv6 capability without a change of the modem on the customer site.

Figure 143 **Dual Stack PPPoE Bridged Gateway Service Example**

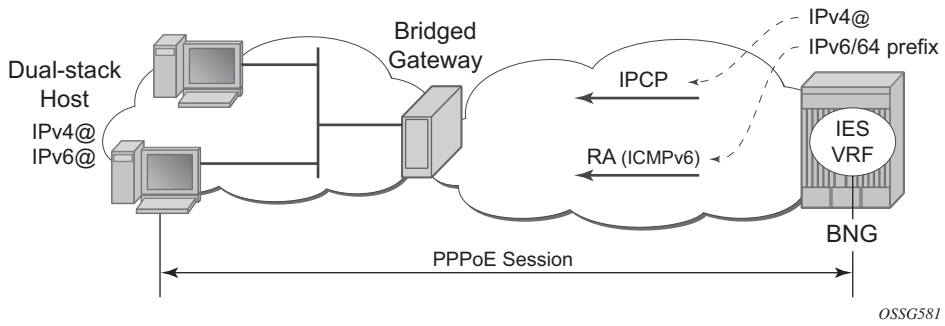

Dual Stack PPPoE Routed Gateway Service

The dual stack PPPoE routed gateway service runs over a dual stack PPPoE session between a dual stack router and BNG. It allows operators using PPPoE in their networks (with either PPPoE to the RG or PPPoA with translation to PPPoE in the DSLAM) to deploy IPv6 services in conjunction with an existing IPv4 service.

Because a routed RG is used, a unique subscriber IPv6 prefix is delegated to the dual stack router for use within the home network. DHCPv6 is used to provide prefix delegation (PD). No WAN IPv6 address assignment is supported in this model. The dual stack router does not perform any NAT for IPv6 traffic.

Figure 144 **Dual Stack PPPoE Routed Gateway Service Example**

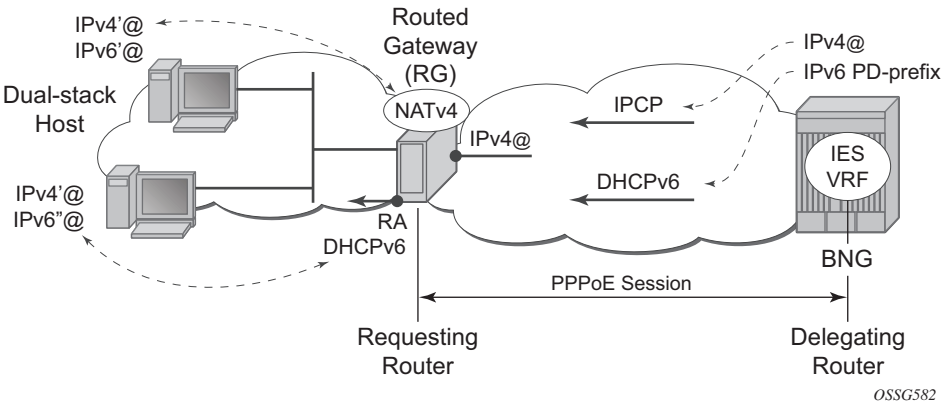

SLAAC

The IPv6 stateless auto configuration (SLAAC) mechanism requires no manual configuration of hosts, minimal configuration of routers, and no additional servers (such as DHCP). The stateless mechanism allows a host to generate its own address using a combination of locally available information and information advertised by routers. Routers advertise /64 prefixes, by an ICMPv6 router advertisement, that identify the subnet(s) associated with a link, while hosts generate a 64-bit “interface identifier” that uniquely identifies an interface on a subnet. An address is formed by combining the two.

DHCPv6

The Dynamic Host Configuration Protocol for IPv6 (DHCPv6) is defined in RFC 3315. The protocol enables DHCPv6 servers to pass configuration parameters such as IPv6 network addresses or DNSv6 addresses to IPv6 nodes.

For further information on DHCPv6, see ESMv6: IPoE Dual Stack Hosts.

Prefix Delegation

Prefix Delegation (PD) is a mechanism for automated delegation of IPv6 prefixes using DHCPv6. A delegating router delegates a long-lived IPv6 prefix to a requesting router. The delegating router does not require knowledge about the topology of the links in the network to which the prefixes will be assigned.

For further information on Prefix Delegation, see ESMv6: IPoE Dual Stack Hosts.

Configuration

ESMv6 for PPPoE is applicable in a routed CO environment. Details of non-specific dual stack configurations like authentication-policies, sla-profile, subscriber-profile, accounting-policies and QoS policies are out of scope for this chapter.

The minimal RADIUS authentication configuration and ESM string configuration is added for completeness.

```
configure
    router
        radius-server
            server "radius-172.16.1.2" address 172.16.1.2 secret vsecret1 create
                accept-coa
            exit
        exit
    exit

configure
    aaa
        radius-server-policy "rsp-1" create
            servers
                router "Base"
                source-address 192.0.2.1
                server 1 name "radius-172.16.1.2"
            exit
        exit
    exit

configure
    subscriber-mgmt
        authentication-policy "auth-1" create
            description "RADIUS authentication policy"
            pppoe-access-method pap-chap
            radius-server-policy "rsp-1"
        exit
    exit

configure
```

```
subscriber-mgmt
    sla-profile "sla-profile-1" create
    exit
    sub-profile "sub-profile-1" create
    exit
    sub-ident-policy "sub-ident-1" create
        sub-profile-map
            use-direct-map-as-default
        exit
        sla-profile-map
            use-direct-map-as-default
        exit
    exit
exit
```

Service

Dual Stack PPPoE for Bridged Gateway

Figure 145 shows the message flow for a dual stack PPPoE host behind a bridged gateway corresponding with the configured service.

Figure 145 **Message Flow for a Dual Stack PPPoE Host**

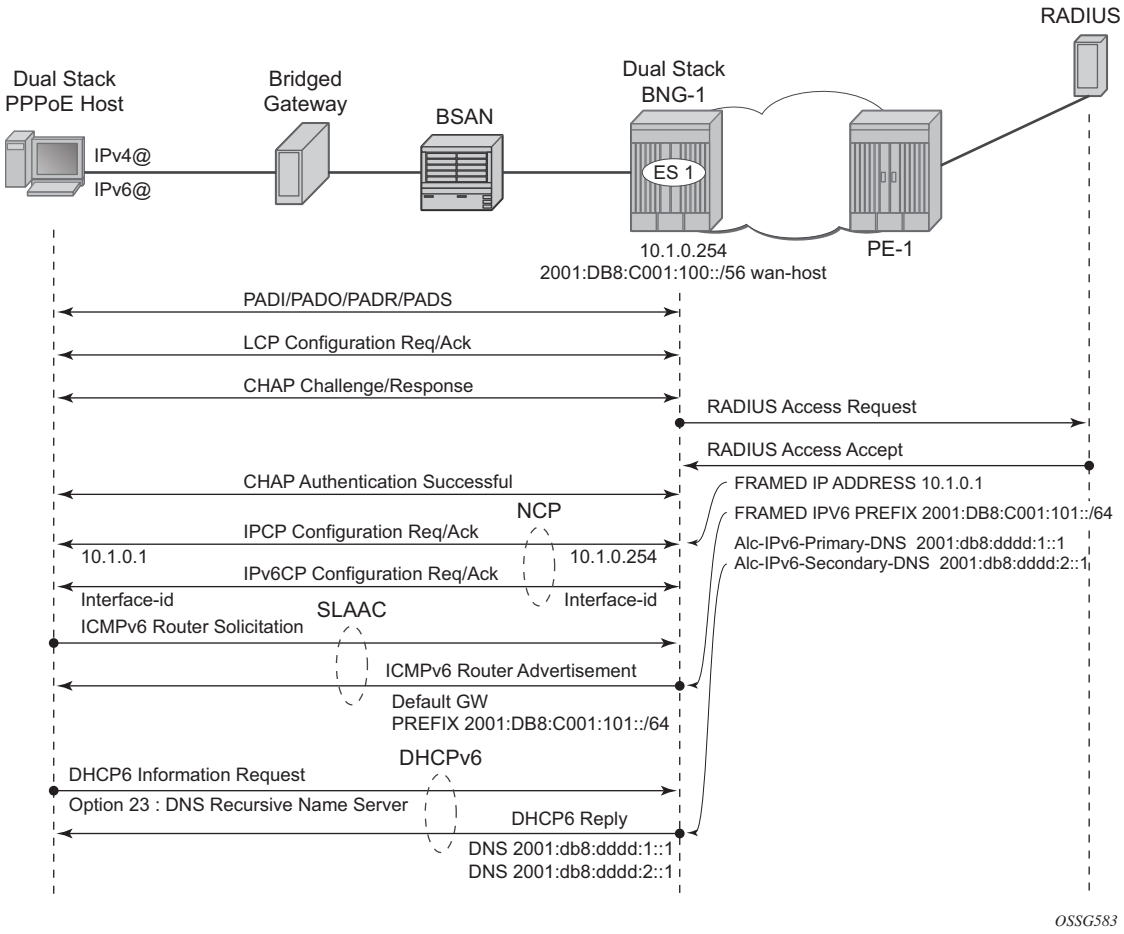

For dual stack PPPoE, the BNG initiates the IPv6 control protocol (IPv6CP) protocol to the client during the session setup phase if the appropriate attributes have been returned by the RADIUS server on authentication. The RADIUS attribute that triggers the setup of a dual stack PPPoE host in bridged mode is *framed-ipv6-prefix* which should contain a /64 prefix for the client. When a PPPoE host has successfully completed the IPv6CP negotiation, the BNG will transmit an RA to the PPPoE host containing the prefix and any other option that is configured. The host can request optional IPv6 DNS server information from the BNG by sending a DHCPv6 information-request.

The following example shows a minimal configuration to enable dual stack subscribers in an IES service context with the ESM IPv6-specific parts in bold.

```
configure
    service
        ies 1 customer 1 create
            subscriber-interface "sub-int-1" create
```

```
                    address 10.1.0.254/16
                    ipv6
                        subscriber-prefixes
                            prefix 2001:db8:c001:100::/56 wan-host
                        exit
                    exit
                    group-interface "group-int-1" create
                        ipv6
                            router-advertisements
                                prefix-options
                                    autonomous
                                exit
                                no shutdown
                            exit
                            dhcp6
                                proxy-server
                                    client-applications ppp
                                    no shutdown
                                exit
                            exit
                        exit
                        authentication-policy "auth-1"
                        sap 1/1/1:1 create
                            sub-sla-mgmt
                                sub-ident-policy "sub-ident-1"
                                multi-sub-sap 10
                                no shutdown
                            exit
                        exit
                        pppoe
                            session-limit 10
                            sap-session-limit 10
                            no shutdown
                        exit
                    exit
                exit
                no shutdown
            exit
        exit
    exit
```

IPv6 subscriber prefixes must be defined in the **subscriber-interface** <*sub-int-name*> **ipv6** *subscriber-prefixes* context.

Three types of prefixes can be configured where **wan-host** is required for the bridged gateway scenario and **pd** is used for the dual stack PPPoE routed gateway scenario.

- wan-host — Prefix from which the IPv6 addresses are assigned (by DHCPv6 IA_NA) for the IPoEv6 routed gateway WAN interface (network facing) or a prefix from which /64 prefixes are assigned for the PPPoE (by RA SLAAC) hosts in the bridged gateway model.
- pd — Prefix from which the IPv6 prefix delegation prefixes are assigned that are to be used by the IPoEv6 or PPPoEv6 routed gateway for allocation in the home network (LAN interfaces).

- pd wan-host (both) — Prefix from which both IPv6 addresses (wan-host) and IPv6 prefix delegation prefixes (pd) can be assigned. This requires that the delegated prefix length is set to 64 bits.

Table 26 and Table 27 provide an overview of the subscriber-prefix parameters that apply and an example of subscriber prefix subnetting for SLAAC.

Table 26 **Subscriber Prefix Parameters**

Subscriber Prefix Type	Prefix Length	DHCPv6 Option	SLAAC	RADIUS AVP	Must be subnetted as
wan-host	/32..63	N/A	yes	[97]Framed-IPv6-Prefix	/64

Table 27 **Subscriber Prefix Subnetting for SLAAC**

Subscriber prefix	Framed-IPv6-Prefix	Hosts	
2001:db8:c001:100::/56	2001:db8:c001:101::/64	pppoev6-host-1	
	2001:db8:c001:102::/64	pppoev6-host-2	
	2001:db8:c001:103::/64	pppoev6-host-3	
	<snip>	<snip>	
	2001:db8:c001:1FF::/64	pppoev6-host-256	
2001:db8:c001:200::/56	2001:db8:c001:201::/64	PPPoEv6-host-257	
	2001:db8:c001:202::/64	PPPoEv6-host-258	
	2001:db8:c001:203::/64	PPPoEv6-host-259	
	<snip>	<snip>	
	2001:db8:c001:2FF::/64	PPPoEv6-host-512	

Dual Stack PPPoE for Routed Gateway

Figure 146 shows the message flow for a dual stack PPPoE host located behind a routed gateway corresponding with the configured service.

Figure 146 **Dual Stack PPPoE for Routed Gateway**

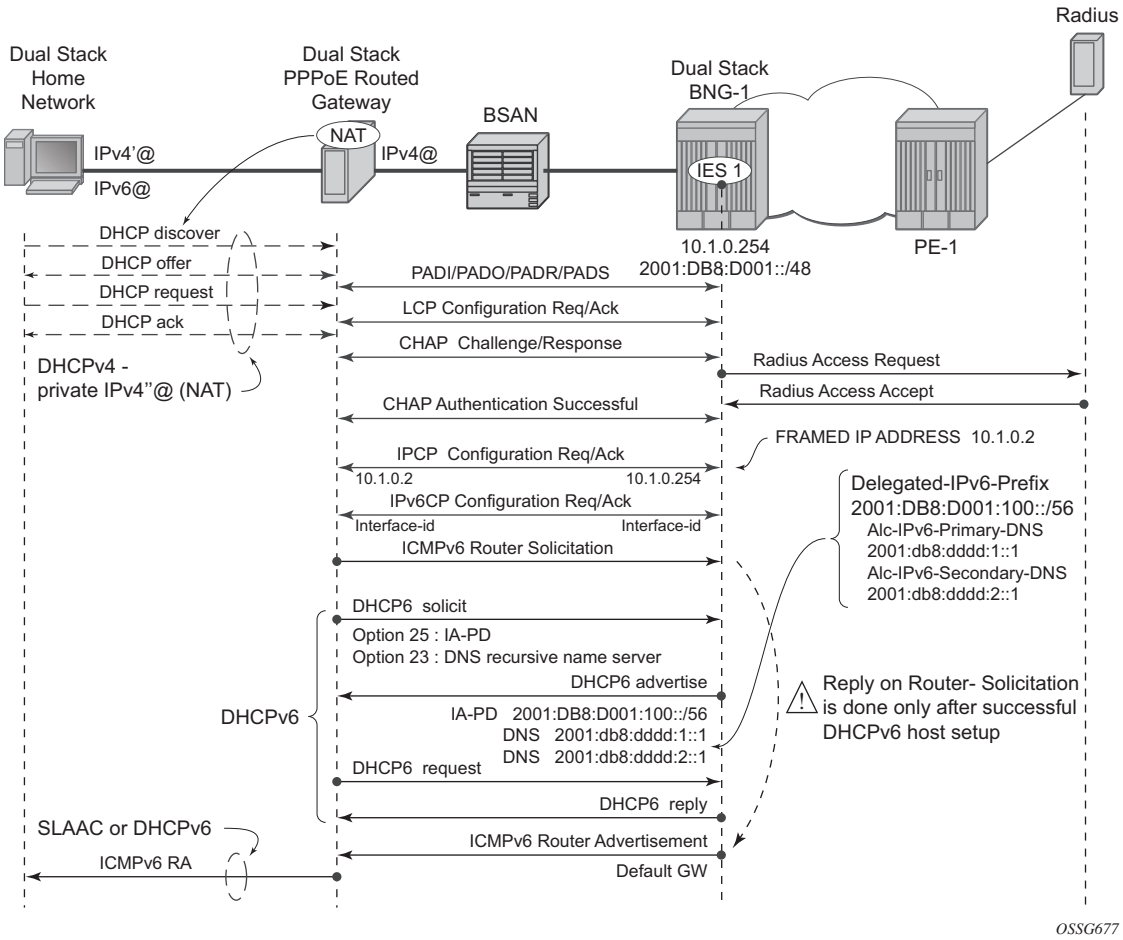

Initially, a PPPoE routed gateway follows the same steps as a dual stack PPPoE host. The BNG receives a prefix from RADIUS (in this case through a *delegated-ipv6-prefix* attribute), which is used as a trigger to initiate the IPv6CP protocol to the client. The prefix that is offered to the client should have the same prefix length as the one configured under the subscriber interface ipv6 context (delegated-prefix length).This length should be between 48 and 64 bits, inclusive.

After the IPv6CP protocol has completed, the client must 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 own downstream interfaces. During the DHCPv6 message exchange, no extra RADIUS request will be made; the information is stored during the initial PPPoE authentication until the client starts DHCPv6. Only after DHCPv6 has completed, the IPv6 subscriber host will be instantiated, and the BNG will start sending RAs if

configured. (It is a mandatory requirement for the BNG to send RAs which makes enabling router-advertisements under the group-level mandatory). The router advertisements do 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 BNG.

```
configure
    service
        ies 1 customer 1 create
            subscriber-interface "sub-int-1" create
                address 10.1.0.254/16
                ipv6
                    delegated-prefix-len 56
                    subscriber-prefixes
                        prefix 2001:db8:d001::/48 pd
                    exit
                exit
                group-interface "group-int-1" create
                    ipv6
                        router-advertisements
                            prefix-options
                                autonomous
                            exit
                            no shutdown
                        exit
                        dhcp6
                            proxy-server
                                client-applications ppp
                                no shutdown
                            exit
                        exit
                    exit
                    authentication-policy "auth-1"
                    sap 1/1/1:1 create
                        sub-sla-mgmt
                            sub-ident-policy "sub-ident-1"
                            multi-sub-sap 10
                            no shutdown
                        exit
                    exit
                    pppoe
                        session-limit 10
                        sap-session-limit 10
                        no shutdown
                    exit
                exit
            exit
            no shutdown
        exit
    exit
exit
```

IPv6 subscriber prefixes must be defined at the **subscriber-interface** <*sub-int-name*> **ipv6** *subscriber-prefixes* context, see Dual Stack PPPoE Bridged Gateway Service.

Subscriber prefixes are subnetted in fixed length subnets that are assigned to subscriber hosts:

- /delegated-prefix-len (/48..64) for p subscriber prefixes

The delegated prefix length is configured in the **subscriber-interface** <*sub-int-name*> **ipv6** context. The recommended value is /56 (default = /64). The configured length applies to all pd subscriber prefixes on a subscriber-interface.

Table 28 and Table 29 provide an overview of the subscriber-prefix parameters that apply and an example of prefix subnetting for delegated-prefix-length /56.

Table 28 **Subscriber-Prefix Parameters**

Subscriber Prefix Type	Prefix Length	DHCPv6 Option	SLAAC	RADIUS AVP	Must be sub netted as
pd	/48..64 *	IA-PD	N/A	[123] Delegated-IPv6-Prefix	/delegated-prefix-len

*Must be smaller than configured delegated prefix length.

Table 29 **Prefix Subnetting for delegated-prefix-length /56**

Subscriber Prefix and /56 delegated-prefix-len	Framed-IPv6-Prefix	Hosts
2001:db8:d001::/48	2001:db8:d001:100::/56	Responsibility Home Gateway (HGW)
		Responsibility HGW
		Responsibility HGW
	2001:db8:d001:200::/56	Responsibility HGW
		Responsibility HGW
		Responsibility HGW
	<snip>	--
	2001:db8:d001:FF00::/56	Responsibility HGW
		Responsibility HGW
		Responsibility HGW
<snip>	--	--

Table 29 **Prefix Subnetting for delegated-prefix-length /56 (Continued)**

Subscriber Prefix and /56 delegated-prefix-len	Framed-IPv6-Prefix	Hosts
2001:db8:d002::/48	2001:db8:d002:100::/56	Responsibility HGW
		Responsibility HGW
		Responsibility HGW
	2001:db8:d002:200::/56	Responsibility HGW
		Responsibility HGW
		Responsibility HGW
	<snip>	--
	2001:db8:d002:FF00::/56	Responsibility HGW
		Responsibility HGW
		Responsibility HGW

RADIUS

The RADIUS authentication policy shown at the beginning of the Configuration section must be applied to the group-interface, and is used for both IPv4 and IPv6.

```
configure
    service
        ies 1 customer 1 create
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    authentication-policy "auth-1"
                exit
            exit
        exit
    exit
exit
```

IPv4 and IPv6 configuration information can come from LUDB or AAA/RADIUS.

Commonly used RADIUS Attribute Value pairs (AVPs) that are applicable for PPPoE IPv6 subscriber hosts are listed in Table 30.

Table 30 **RADIUS AVPs**

RADIUS AVP	Type	Purpose
Framed-IPv6-Prefix [97]	ipv6prefix	Maps to SLAAC (RFC 4862) /64 Prefix-information in ICMPv6 RA.
Delegated-IPv6-Prefix [123]	ipv6prefix	Maps to IA_PD for prefix delegation (RFC 3633) in DHCPv6
Alc-Ipv6-Primary-Dns [26-6527-105]	ipv6addr	Maps to DNS recursive name server option (RFC 3646) in DHCPv6
Alc-Ipv6-Secondary-Dns [26-6527-106]	ipv6addr	Maps to DNS recursive name server option (RFC 3646) in DHCPv6

Dual Stack PPPoE for Bridged Gateway

The following shows a sample of a FreeRADIUS user record to authenticate a dual stack PPPoE subscriber for a bridged gateway:

```
bridged@domain1 Cleartext-Password := "letmein"
    Framed-IP-Address = 10.1.0.1,
    Framed-IP-Netmask = 255.255.255.0,
    Alc-Subsc-ID-Str = "%{User-name}",
    Alc-Subsc-Prof-Str = "sub-profile-1",
    Alc-SLA-Prof-Str = "sla-profile-1",
    Framed-IPv6-Prefix = "2001:db8:c001:0101::/64",
    Alc-IPv6-Primary-DNS = "2001:db8:dddd:1::1",
    Alc-IPv6-Secondary-DNS = "2001:db8:dddd:2::1",
```

Dual Stack PPPoE for Routed Gateway

The following shows a sample of a FreeRADIUS user record to authenticate a dual stack PPPoE subscriber for a routed gateway:

```
routed@domain1 Cleartext-Password := "letmein"
    Framed-IP-Address = 10.1.0.2,
    Framed-IP-Netmask = 255.255.255.0,
    Alc-Subsc-ID-Str = "%{User-name}",
    Alc-Subsc-Prof-Str = "sub-profile-1",
    Alc-SLA-Prof-Str = "sla-profile-1",
    Delegated-IPv6-Prefix = "2001:db8:d001:0100::/56",
    Alc-IPv6-Primary-DNS = "2001:db8:dddd:1::1",
    Alc-IPv6-Secondary-DNS = "2001:db8:dddd:2::1",
```

A RADIUS user's configuration with multiple delegated-ipv6-prefixes for the same dual stack PPPoE host will result in a single DHCPv6 advertise message sent by the BNG with a single IA_PD option and single IA-Prefix. The other RADIUS configured delegated-IPv6-prefixes are silently dropped by the BNG.

Router Advertisements

ICMPv6 router advertisements have two major functions.

- Default router function for hosts
- Address auto-configuration for hosts aka SLAAC

Unsolicited RA must explicitly be enabled on a group interface (default shutdown) and are refreshed with a pseudo random timer. The boundaries of this random timer are configurable with the min-advertisement parameter (minimum with default set to 900s) and max-advertisement (maximum with default set to 1800s).

```
configure
    service
        ies 1 customer 1 create
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    ipv6
                        router-advertisements
                            max-advertisement 1800  # default 30 min
                            min-advertisement 900   # default 15 min
                            no shutdown
                        exit
                    exit
                exit
            exit
        exit
    exit
```

The **router-advertisements router-lifetime** parameter (default 4500 sec) specifies how long the host is allowed to use the originator of the RA as default gateway. This timer is configurable between 2700 and 9000 seconds.

Configuring a **router-advertisements router-life** timer smaller than the **router-advertisements min-advertisement** timer results in a dual stack PPPoE host without a default gateway.

```
configure
    service
        ies 1 customer 1 create
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    ipv6
                        router-advertisements
```

```
                        router-lifetime 4500
                        no shutdown
                    exit
                exit
            exit
        exit
    exit
```

The following **prefix-options autonomous** parameter specifies whether or not offered RADIUS IPv6 prefix can be used for stateless address configuration (SLAAC). The **prefix-options lifetime** parameter defines how long the host is allowed to use this prefix. Configuring a **prefix-option valid-lifetime** smaller than the **router-advertisements min-advertisement** timer results in host traffic being sourced with the link-local address instead of global unique IPv6 address.

```
configure
    service
        ies 1 customer 1 create
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    ipv6
                        router-advertisements
                            prefix-options
                                autonomous                # required for SLAAC
                                on-link
                                preferred-lifetime 3600   # default 1 hour
                                valid-lifetime 86400      # default 24 hours
                            exit
                            no shutdown
                        exit
                    exit
                exit
            exit
        exit
    exit
exit
```

The following is a snapshot from an ICMPv6 RA message with default timer settings with a focus on the SLAAC function.

```
Internet Control Message Protocol v6
    Type: 134 (Router advertisement)

    --- snip ---

    ICMPv6 Option (Prefix information)
        Type: Prefix information (3)
        Length: 32
        Prefix length: 64
        Flags: 0x40
            1... .... = on link
            .1.. .... = Auto                    # Auto-Configuration flag
            ..0. .... = Not router address
            ...0 .... = Not site prefix
```

```
Valid lifetime: 86400                  # Default value 24 hour
Preferred lifetime: 3600               # Default value  1 hour
Prefix: 2001:DB8:C001:101::            # SLAAC prefix
```

SLAAC-related parameters are listed in Table 31.

Table 31 **SLAAC-Related Parameters**

Parameter	Description (RFC-4861)	Value Range (Default)
prefix-options: autonomous	Autonomous address-configuration flag. When set indicates that this prefix can be used for stateless address auto configuration (SLAAC)	(no)
prefix-options: preferred-lifetime	The length of time in seconds that the addresses generated from the prefix through stateless address auto configuration (SLAAC) remains preferred.	0..4294967295 s (3600s) 1hour
prefix-options: valid-lifetime	The length of time in seconds that the prefix is valid for the purpose of on-link determination.	0..4294967295 s (86400s) 24hours

Router advertisements parameters common to PPPoEv6 and IPoEv6 are listed and explained in ESMv6: IPoE Dual Stack Hosts.

For dual stack PPPoE hosts, the default values, as shown in the following output, can be used. Timer values equal to zero (reachable-time and retransmit-time) causes the host to use its own timers for that function. The reachable-time is used by the host for Neighbor_Unreachable_Detection (NUD) whereas the retransmit-time is used by the host for Duplicate_Address_Detection (DAD). DAD is normally only performed by dual stack IPoE hosts and not by dual stack PPPoE hosts.

```
configure
    service
        ies 1 customer 1 create
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    ipv6
                        router-advertisements
                            current-hop-limit 64
                            dns-options
                                no include-dns
                                rdnss-lifetime 3600
                            exit
                            no force-mcast
                            no managed-configuration
                            no mtu
                            no other-stateful-configuration
                            reachable-time 0
                            retransmit-time 0
                            no shutdown
```

```
                    exit
                exit
            exit
        exit
    exit
```

DHCPv6 Proxy Server

Dual Stack PPPoE for Bridged Gateway

Dual stack PPPoE hosts using SLAAC for address assignment do not require DHCPv6. SR OS supports DNSv6 information through the RA DNS Option (RFC 5006, *IPv6 Router Advertisement Option for DNS Configuration*), and can also be configured to use DHCPv6 information requests and replies to retrieve the DNSv6 information. This requires the DHCPv6 proxy server to be enabled (the default is **shutdown**) and PPPoE defined as client-application (default=dhcp only). No lease state is kept for this DNSv6 information and therefore it is known as stateless DHCPv6.

```
configure
    service
        ies 1 customer 1 create
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    ipv6
                        dhcp6
                            proxy-server
                                client-applications ppp
                                no shutdown
                            exit
                        exit
                    exit
                exit
            exit
        exit
```

Dual Stack PPPoE for Routed Gateway

An IPv6 PPPoE routed gateway initiates, after successful IPv6CP negotiation, a DHCPv6 session to request its configuration data (IPv6 PD prefixes, DNS servers). A DHCPv6 proxy server in the BNG maintains the DHCPv6 session with the IPv6 PPPoE subscriber host. The DHCPv6 proxy server must be enabled (the default is **shutdown**) and PPPoE defined as client-application (default=dhcp only).

```
configure
    service
        ies 1 customer 1 create
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    ipv6
                        dhcp6
                            proxy-server
                                server-id duid-ll
                                renew-timer min 30        # default
                                rebind-timer min 48       # default
                                valid-lifetime days 1     # default
                                preferred-lifetime hrs 1  # default
                                client-applications ppp
                                no shutdown
                            exit
                        exit
                    exit
                exit
```

A number of timers associated with IPv6 addresses and IPv6 prefixes within DHCPv6 identity associations can be configured in the DHCPv6 proxy server context. These timers are valid for IPoEv6 and PPPoEv6 sessions and are listed and further explained in ESMv6: IPoE Dual Stack Hosts.

There is never RADIUS re-authentication for dual stack PPPoE routed gateways on DHCPv6 renewals as indicated in Figure 147.

Figure 147 **DHCPv6 Renewals**

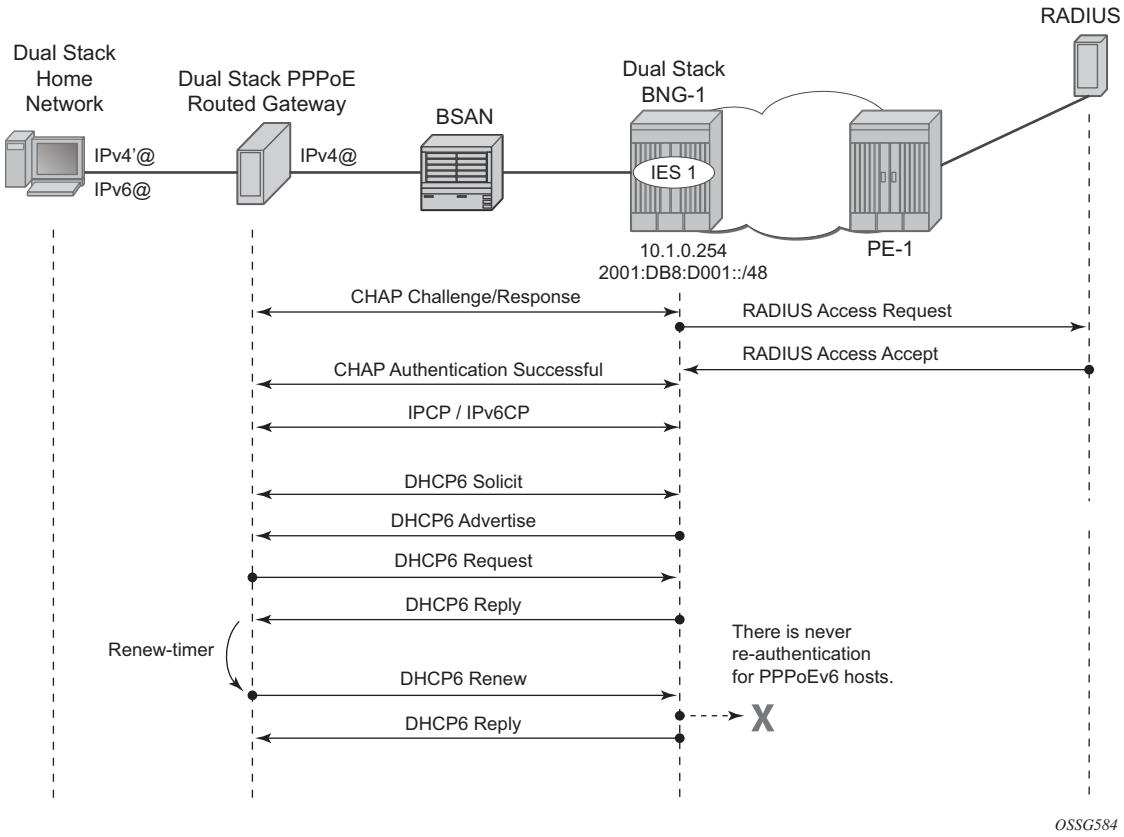

DHCPv6 Lease State

The DHCPv6 lease state table keeps track of the DHCPv6 host states. The DHCP lease information for a specific host is extracted from the DHCPv6 reply message in case of DHCPv6. Stateful (with lease state) DHCPv6 is applicable for dual stack PPPoE on routed gateway where Stateless (without lease state) DHCPv6 is optional and applicable for dual stack PPPoE on bridged gateways.

For more information on DHCPV6 lease states, see ESMv6: IPoE Dual Stack Hosts.

Operation

Dual Stack PPPoE for Bridged Gateway

A PPPoEv6 dual stack subscriber scenario for a bridged home gateway consumes two subscriber host entries sharing a common subscriber.

- IPv4 host-addressing by IPCP
- IPv6 wan-host addressing by SLAAC

```
*A:BNG# show service active-subscribers subscriber bridged@domain1"bridged@domain1"
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber bridged@domain1 (sub-profile-1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:1/1/1:1 - sla:sla-profile-1
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
10.1.0.1
                00:0c:29:00:00:11    PPP 1          IPCP         1          Y
2001:db8:c001:101::/64
                00:0c:29:00:00:11    PPP 1          SLAAC        1          Y
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
*A:BNG#
```

The **hierarchy** parameter for active-subscribers gives a top level down overview for this subscriber.

```
*A:BNG# show service active-subscribers hierarchy subscriber "bridged@domain1"

===============================================================================
Active Subscribers Hierarchy
===============================================================================
Hierarchy
-------------------------------------------------------------------------------
-- bridged@domain1 (sub-profile-1)
   |
   +-- sap:1/1/1:1 - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:0c:29:00:00:11 - sid:1 - svc:1
           |
           |-- 10.1.0.1 - IPCP
           |
           +-- 2001:db8:c001:101::/64 - SLAAC

===============================================================================
```

```
*A:BNG#
```

IPCP and IPv6CP are in an opened state for the dual stack PPPoE session and their origin is RADIUS, as shown below.

```
*A:BNG# show service id 1 pppoe session ip-address 10.1.0.1 detail

===============================================================================
PPPoE sessions for svc-id 1
===============================================================================
Sap Id               Mac Address       Sid    Up Time         Type
    IP/L2TP-Id/Interface-Id                                      MC-Stdby
-------------------------------------------------------------------------------
1/1/1:1              00:0c:29:00:00:11 1      0d 00:00:46     local
    10.1.0.1
    02:0C:29:FF:FE:00:00:11

LCP State            : Opened
IPCP State           : Opened
IPv6CP State         : Opened
PPP MTU              : 1492
PPP Auth-Protocol    : CHAP
PPP User-Name        : bridged@domain1

Subscriber-interface : sub-int-1
Group-interface      : group-int-1

IP Origin            : radius
DNS Origin           : none
NBNS Origin          : none

Subscriber           : "bridged@domain1"
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-1"

--- snipped ---

IPv6 Prefix          : 2001:db8:c001:101::/64
IPv6 Prefix Origin   : radius
IPv6 Prefix Pool     : ""
IPv6 Del.Pfx.        : N/A
IPv6 Del.Pfx. Origin : none
IPv6 Del.Pfx. Pool   : ""
IPv6 Address         : N/A
IPv6 Address Origin  : none
IPv6 Address Pool    : ""
Primary IPv6 DNS     : 2001:db8:dddd:1::1
Secondary IPv6 DNS   : 2001:db8:dddd:2::1

--- snipped ---

-------------------------------------------------------------------------------
Number of sessions   : 1
===============================================================================
*A:BNG#
```

The IPv6 routing table for dual stack hosts is displayed using the **protocol** keyword **sub-mgmt**.

```
*A:BNG# show router route-table ipv6 protocol sub-mgmt

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age         Pref
      Next Hop[Interface Name]                                     Metric
-------------------------------------------------------------------------------
2001:db8:c001:101::/64                          Remote  Sub Mgmt  00h00m57s  0
       [group-int-1]                                                 0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
*A:BNG#
```

DNSv6

DNSv6 information, in a dual stack PPPoE bridged gateway model, is optionally retrieved through stateless DHCPv6 information requests. Debugging is done through debug commands or/and observation by statistics counters.

```
100 2016/07/08 14:16:40.01 CEST MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Incoming DHCP6 Msg : INFO_REQUEST (11)
   on itf group-int-1
     Trans Id : 0xcef3f0
     Option : CLIENTID (1), Length : 14
       LLT : HwTyp=0001,T=322878930,LL=000c29c851ca
       00010001133ebdd2000c29c851ca
     Option : ELAPSED_TIME (8), Length : 2
       Time : 100 seconds
     Option : ORO (6), Length : 4
       Requested Option : DNS_NAME_SRVR (23)

101 2016/07/08 14:16:40.02 CEST MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Outgoing DHCP6 Msg : REPLY (7)
   to itf group-int-1
     Trans Id : 0xcef3f0
     Option : SERVERID (2), Length : 10
       LL  : HwTyp=0001,LL=24b1ff000000
       0003000124b1ff000000
     Option : CLIENTID (1), Length : 14
       LLT : HwTyp=0001,T=322878930,LL=000c29c851ca
       00010001133ebdd2000c29c851ca
     Option : DNS_NAME_SRVR (23), Length : 32
       Server : 2001:db8:dddd:1::1
       Server : 2001:db8:dddd:2::1
```

DHCPv6 statistics

```
*A:BNG# 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

 --- snipped ---

38 Packet dropped by DHCP filter                                  0
39 Packet dropped because authentication failed                   0
===============================================================================
*A:BNG#
```

To clear the statistics use following command:

```
*A:BNG# clear router dhcp6 statistics
```

Entries, for dual stack PPPoE subscribers, in the IPv4 ARP and/or IPv6 neighbor cache table are counted as internal entries and are shown from the **summary** parameter.

```
*A:BNG# show router arp summary

===============================================================
ARP Table Summary (Router: Base)
===============================================================
Local ARP Entries    : 3
Static ARP Entries   : 0
Dynamic ARP Entries  : 1
Managed ARP Entries  : 0
```

```
Internal ARP Entries : 1
BGP-EVPN ARP Entries : 0
-------------------------------------------------------
No. of ARP Entries   : 5
=======================================================
*A:BNG#

*A:BNG# show router neighbor summary

===============================================================
Neighbor Table Summary (Router: Base)
===============================================================
Static Nbr Entries      : 0
Dynamic Nbr Entries     : 0
Managed Nbr Entries     : 0
Internal Nbr Entries    : 1
Evpn Nbr Entries        : 0
---------------------------------------------------------------
No. of Neighbor Entries : 1
===============================================================
*A:BNG#
```

Dual Stack PPPoE for Routed Gateway

A PPPoEv6 dual stack subscriber scenario for a routed CPE consumes two subscriber host entries sharing a common subscriber.

- IPv4 host addressing through IPCP
- IPv6 pd addressing through DHCPv6

```
*A:BNG# show service active-subscribers subscriber routed@domain1"routed@domain1"

===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber routed@domain1 (sub-profile-1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:1/1/1:1 - sla:sla-profile-1
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
10.1.0.2
                00:0c:29:00:00:12    PPP 1          IPCP         1          Y
2001:db8:d001:100::/56
                00:0c:29:00:00:12    PPP 1          DHCP6-PD     1          Y
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
*A:BNG#
```

The hierarchy parameter for active-subscribers gives a top level down overview for this subscriber.

```
*A:BNG# show service active-
subscribers hierarchy subscriber routed@domain1"routed@domain1"

===============================================================================
Active Subscribers Hierarchy
===============================================================================
Hierarchy
-------------------------------------------------------------------------------
-- routed@domain1 (sub-profile-1)
   |
   +-- sap:1/1/1:1 - sla:sla-profile-1
       |
       +-- PPP-session - mac:00:0c:29:00:00:12 - sid:1 - svc:1
           |
           |-- 10.1.0.2 - IPCP
           |
           +-- 2001:db8:d001:100::/56 - DHCP6-PD

===============================================================================
*A:BNG#
```

IPCP and IPv6CP are in an opened state for the dual stack PPPoE session and their origin is RADIUS, as shown below.

```
*A:BNG# show service id 1 pppoe session ip-address 10.1.0.2 detail

===============================================================================
PPPoE sessions for svc-id 1
===============================================================================
Sap Id               Mac Address       Sid    Up Time         Type
    IP/L2TP-Id/Interface-Id                                      MC-Stdby
-------------------------------------------------------------------------------
1/1/1:1              00:0c:29:00:00:12 1      0d 00:00:43     local
    10.1.0.2
    02:0C:29:FF:FE:00:00:12

LCP State            : Opened
IPCP State           : Opened
IPv6CP State         : Opened
PPP MTU              : 1492
PPP Auth-Protocol    : CHAP
PPP User-Name        : routed@domain1

Subscriber-interface : sub-int-1
Group-interface      : group-int-1

IP Origin            : radius
DNS Origin           : none
NBNS Origin          : none

Subscriber           : "routed@domain1"
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-1"
```

```
--- snipped ---

IPv6 Prefix          : N/A
IPv6 Prefix Origin   : none
IPv6 Prefix Pool     : ""
IPv6 Del.Pfx.        : 2001:db8:d001:100::/56
IPv6 Del.Pfx. Origin : radius
IPv6 Del.Pfx. Pool   : ""
IPv6 Address         : N/A
IPv6 Address Origin  : none
IPv6 Address Pool    : ""
Primary IPv6 DNS     : 2001:db8:dddd:1::1
Secondary IPv6 DNS   : 2001:db8:dddd:2::1

--- snipped ---

-------------------------------------------------------------------------------
Number of sessions   : 1
===============================================================================
*A:BNG#
```

The IPv6 routing table for dual stack hosts is displayed using the **protocol** keyword **sub-mgmt**.

```
*A:BNG# show router route-table ipv6 protocol sub-mgmt

===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age         Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:db8:d001:100::/56                        Remote  Sub Mgmt  00h00m51s  0
       [group-int-1]                                                 0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
*A:BNG#
```

DNSv6

DNSv6 information, in a dual stack PPPoE routed gateway model, is optionally retrieved by stateful DHCPv6 information requests. Debugging is done through debug commands or/and observation by statistics counters.

```
14 2016/07/08 09:48:36.62 CEST MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Incoming DHCP6 Msg : SOLICIT (1)
   on itf group-int-1
     Trans Id : 0x6c8a21
```

```
      Option : CLIENTID (1), Length : 14
        LLT : HwTyp=0001,T=521204519,LL=000c29000012
        000100011f10f327000c29000012
      Option : IA_PD (25), Length : 12
        IAID : 1
        Time1: 0 seconds
        Time2: 0 seconds
      Option : ORO (6), Length : 2
        Requested Option : DNS_NAME_SRVR (23)
"

15 2016/07/08 09:48:36.62 CEST MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Outgoing DHCP6 Msg : ADVERTISE (2)
   to itf group-int-1
     Trans Id : 0x6c8a21
     Option : SERVERID (2), Length : 10
       LL  : HwTyp=0001,LL=02c9ff000000
       0003000102c9ff000000
     Option : CLIENTID (1), Length : 14
       LLT : HwTyp=0001,T=521204519,LL=000c29000012
       000100011f10f327000c29000012
     Option : DNS_NAME_SRVR (23), Length : 32
       Server : 2001:db8:dddd:1::1
       Server : 2001:db8:dddd:2::1
     Option : IA_PD (25), Length : 41
       IAID : 1
       Time1: 1800 seconds
       Time2: 2880 seconds
       Option : IAPREFIX (26), Length : 25
         Prefix : 2001:db8:d001:100::/56
         Preferred Lifetime : 3600 seconds
         Valid Lifetime     : 86400 seconds
"

16 2016/07/08 09:48:36.63 CEST MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Incoming DHCP6 Msg : REQUEST (3)
   on itf group-int-1
     Trans Id : 0x81edfc
     Option : CLIENTID (1), Length : 14
       LLT : HwTyp=0001,T=521204519,LL=000c29000012
       000100011f10f327000c29000012
     Option : SERVERID (2), Length : 10
       LL  : HwTyp=0001,LL=02c9ff000000
       0003000102c9ff000000
     Option : IA_PD (25), Length : 12
       IAID : 1
       Time1: 0 seconds
       Time2: 0 seconds
     Option : ORO (6), Length : 2
       Requested Option : DNS_NAME_SRVR (23)
"

17 2016/07/08 09:48:36.63 CEST MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Outgoing DHCP6 Msg : REPLY (7)
   to itf group-int-1
     Trans Id : 0x81edfc
```

```
      Option : SERVERID (2), Length : 10
        LL  : HwTyp=0001,LL=02c9ff000000
        0003000102c9ff000000
      Option : CLIENTID (1), Length : 14
        LLT : HwTyp=0001,T=521204519,LL=000c29000012
        000100011f10f327000c29000012
      Option : DNS_NAME_SRVR (23), Length : 32
        Server : 2001:db8:dddd:1::1
        Server : 2001:db8:dddd:2::1
      Option : IA_PD (25), Length : 41
        IAID : 1
        Time1: 1800 seconds
        Time2: 2880 seconds
        Option : IAPREFIX (26), Length : 25
          Prefix : 2001:db8:d001:100::/56
          Preferred Lifetime : 3600 seconds
          Valid Lifetime     : 86400 seconds
"
```

Use the following command to display the DHCPv6 statistics.

```
*A:BNG# show router dhcp6 statistics

===============================================================================
DHCP6 statistics (Router: Base)
===============================================================================
Msg-type                        Rx              Tx              Dropped
-------------------------------------------------------------------------------
1 SOLICIT                       1               0               0
2 ADVERTISE                     0               1               0
3 REQUEST                       1               0               0
4 CONFIRM                       0               0               0
5 RENEW                         0               0               0
6 REBIND                        0               0               0
7 REPLY                         0               1               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
--- snipped ---
38 Packet dropped by DHCP filter                                  0
39 Packet dropped because authentication failed                   0
===============================================================================
*A:BNG#
```

Use the following command to clear the DHCPv6 statistics.

```
A:BNG-1# clear router dhcp6 statistics
```

Debugging

Following tools are available for troubleshooting PPPoE dual stack scenarios.

- system log (log-id 99)
- debugging aids
- protocol statistics

Log-id 99 is the default system log. Use appropriate filtering to reduce the output if needed.

```
*A:BNG# show log log-id 99
```

Following debug configuration is useful for troobleshooting PPPoE, RADIUS, DHCPv6 and ICMPv6.

```
debug
    service
        id 1
            ppp
                packet
                    mode egr-ingr-and-dropped
                    detail-level high
                    discovery
                    ppp
                    dhcp-client
                exit
            exit
        exit
    exit
exit
```

```
debug
    router
        radius
            packet-type authentication accounting coa
            detail-level high
        exit
    exit
exit

debug
```

```
    router "Base"
        ip
            dhcp6
                mode egr-ingr-and-dropped
                detail-level high
            exit
            icmp6
        exit
    exit
exit
```

Use the following commands for showing protocol related statistics.

```
show router dhcp6 statistics
show service id 1 pppoe session statistics
```

Advanced Topics

RADIUS COA

For dual stack PPPoE subscriber hosts, RADIUS-triggered mid-session change or/and session terminations identify the subscriber host to be changed by the same prefix that was originally returned from RADIUS or by the host-session-id (If RADIUS accounting host-accounting is enabled and the accounting session-id format equals number). Changing either the IPv4 or IPv6 information will result in both the v4 and v6 subscriber hosts being modified. Further elaboration on accounting is out of scope in this document.

IPv6CP Interface ID

IPv6CP negotiates, unlike ipv4-addresses in IPv4CP, only interface-ids (interface-id: the last 64 bits of an IPv6 address is the interface identifier that is unique to the 64-bit prefix of the IPv6 address and is usually derived from the link-layer or MAC address).

Dual stack PPPoE subscribers and the BNG exchange their interface-ids during the NCP phase. For ESM subscriber-interfaces on the BNG the interface-id is derived from the chassis-mac address.

- The BNG will nack the PPPoE host's IPv6CP configuration request if the dual stack PPPoE host negotiates an interface-id equal zero or an interface-id equal to the BNG interface ID. In that scenario, the BNG offers in the IPv6CP nack message a suitable interface ID, see Table 32.
- The BNG terminates the session if the dual stack PPPoE hosts nacks its IPv6CP configuration request and offers something else to the BNG.

Table 32 **IPv6CP Nack Message Format**

1	2	3	4	5	6	7	8
SAP ID				Last 2 bytes MAC host		Session ID	

Conclusion

This chapter provides configuration and troubleshooting commands for dual stack PPPoE subscribers on bridged or routed gateways. SLAAC is used as IPv6 address assignment for bridged gateway scenarios and stateful DHCPv6 prefix-delegation is used for address assignment for routed gateway scenarios. No RG WAN IPv6 address assignment is supported in this latter model.

DNSv6 addressing on a bridged gateway is retrieved by stateless DHCPv6 (information request and reply).

Establishing a Diameter Peering Session

This chapter provides information about establishing a diameter peering session.

Topics in this chapter include:

Applicability

This example is applicable to all 7750 SR/SR-c and 7450 ESS chassis.

The configuration was tested on release 12.0.R3.

Overview

Diameter is an Authentication, Authorization and Accounting (AAA) protocol that has been defined by the IETF in RFC 3588, *Diameter Base Protocol*, as a replacement for other AAA protocols like TACACS and RADIUS. While wireline access networks are largely based on RADIUS for subscriber authentication, authorization, and accounting, it was decided by 3rd Generation Partnership Project (3GPP) that wireless access networks will be largely based on Diameter. Over time, operators are looking to converge both types of networks, and one of the aspects of this is to replace RADIUS in wireline access networks by Diameter.

Diameter is based on three layers: the transport layer, the Diameter base protocol layer and the Diameter applications as shown in Figure 148.

Figure 148 **Diameter Protocol Stack**

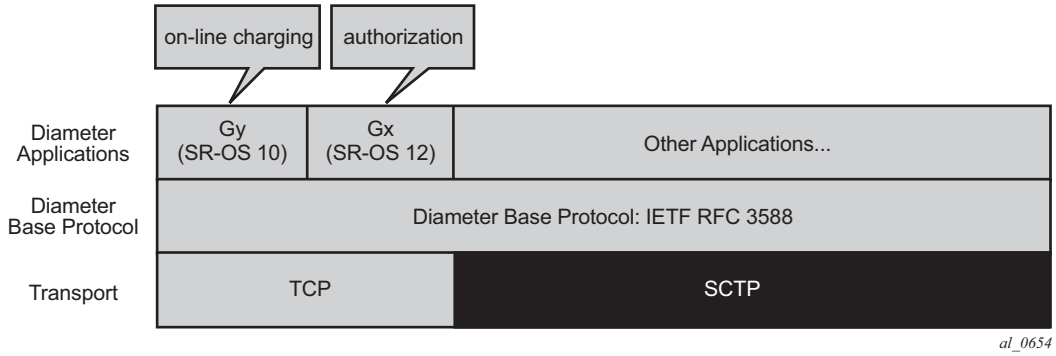

The bottom layer is the transport layer and can be either TCP or SCTP, but only TCP is supported by SR OS. The Diameter base protocol implementation is based on RFC 3588. The top layer contains the Diameter applications. SR OS 10.0 introduced the Gy Diameter application for on-line charging, and SR OS 12.0 introduces the Gx Diameter application for authorization.

Configuration

The Diameter base protocol and the Diameter applications have to be configured separately, where the Diameter base protocol has to be configured first, and the Diameter applications next. The transport layer configuration is part of the Diameter base protocol layer. This example only describes the Diameter base protocol configuration.

The diameter-peer-policy configuration resides in the aaa context and contains the full Diameter base protocol configuration. An example diameter-peer-policy configuration is shown below.

```
configure aaa
    diameter-peer-policy "DSC.26.206" create
        applications gx
        origin-host "wlangw-2.SRrealm"
        origin-realm "SRrealm"
        router 10000
        source-address 10.23.0.130
        peer "DSC.26.206" create
            address 10.40.11.2
            destination-realm "Tc3eRealm"
            no shutdown
        exit
```

The diameter-peer-policy is identified by name (**DSC.26.206** in the example above). This name is used by the Diameter application configuration. The Diameter application making use of this Diameter peer policy must be specified (for instance Gx), and so is the Diameter identity. The Diameter identity consists of 2 parts: the Diameter host name and the Diameter realm, **wlangw-2.SRrealm** and **SRrealm**, respectively, in the example above. Configuration of the Diameter application making use of the Diameter peer is required such that the Diameter connection including the capability exchange, which negotiates the Diameter application with the peer, can already start without configuration of the specific Diameter application.

By default the system originates the Diameter peering session from the Base router but a different routing context (a VPRN or the management routing context used for out-of-band management) can be used (VPRN **10000** in the example). Also a source address (belonging to an IP interface in the configured routing context) can be specified (**10.23.0.130** in the example), but when no source address is specified an address is selected automatically. As such, best practice is to explicitly configure the source address.

One or more peers can be configured, with a maximum of five peers. In case more than one peer is configured, these peers can provide redundancy when supported by the Diameter application making use of the Diameter peer policy. Each peer is identified by name (which could be the same as the Diameter peer policy as in the example above) and has an IP address and a destination Diameter realm (**10.40.11.2** and **Tc3eRealm**, respectively, in the example). The IP address is that of the device terminating the TCP session, which either is the final destination or a Diameter Routing Agent (DRA) (intermediate destination). The destination realm is typically the realm of the final destination. Optionally the destination host can be configured. If the destination host is not configured, then it will be dynamically learned from the received Diameter application messages (learned from the received origin-host). This means that if it is not configured, the first Diameter application message does not contain a destination-host, only a destination-realm, but all subsequent messages will include the learned destination-host. If destination-host is configured, it will be included in the first Diameter application message. The destination host is configured per peer using following command:

```
*A:BNG-1# configure aaa diameter-peer-
policy "DSC.26.206" peer "DSC.26.206" destination-host
  - destination-host <destination-host-string>
  - no destination-host

 <destination-host-*> : [80 chars max]
```

Configuration of the origin-host, origin-realm, destination-realm and at least 1 peer is mandatory. These attributes do not have default values and are needed before a peer can be put in a **no shutdown** state. Doing a **no shutdown** of a peer fails in case any of these attributes are not configured, for instance:

```
*A:BNG-1>config>aaa>diameter-peer-plcy>peer$ no shutdown
MINOR: DIAM #1205 Origin-host is not configured yet

*A:BNG-1>config>aaa>diam-peer-plcy# peer "test" no shutdown
MINOR: DIAM #1206 Origin-realm is not configured yet

*A:BNG-1>config>aaa>diam-peer-plcy# peer "test" no shutdown
MINOR: DIAM #1208 Destination-realm is not configured yet
```

When doing a **no shutdown** of the peer, the system tries to establish the TCP session. Once the TCP session is up, the system starts the Diameter capability negotiation using the configured attributes: the Diameter identity is advertised together with the configured Diameter applications. An example of a capability negotiation is examined in detail in the troubleshooting section. All Diameter messages are sent with a DSCP set to AF41. The DSCP value cannot be changed.

The status of the Diameter peers can be verified as follows:

```
*A:BNG-1# show aaa diameter-peer-policy "DSC.26.206"
===============================================================================
Diameter Peer Policy : DSC.26.206
===============================================================================
Last Mgmt Change      : 03/18/2014 17:50:50
-------------------------------------------------------------------------------
Diameter Base Values (config)
-------------------------------------------------------------------------------
Origin Host           : wlangw-2.SRrealm
Origin Realm          : SRrealm
Connection Timer      : 30 (default)       Source Address     : 10.23.0.130
Transaction Timer     : 30 (default)       Router             : 10000
Watchdog Timer        : 30 (default)
Vendor Support        : 3GPP (default)
-------------------------------------------------------------------------------
Peer Name            Oper  PSM State     Susp  Cooldown  Pref  Order  Pri/Sec
-------------------------------------------------------------------------------
DSC.26.206           Yes   I-Open        No    -         50    1      Primary
===============================================================================
```

An important information is the **PSM State** of each peer. The state **I-Open** indicates that the peer is up and running. The full Peer State Machine (PSM) is described in RFC 3588.

More advanced configuration can be done as well. Timers can be configured: connection timer, transaction timer and watchdog timer:

- The connection timer is called the Tc timer in RFC 3588. This timer controls the frequency at which a new connection is attempted to be established. The default value is 30 seconds as recommended by RFC 3588.

- The transaction timer is started each time a request is sent to the peer and indicates the time the system waits for an answer before resending the request to one of the other configured peers. In case the Diameter request is retransmitted to another peer, the T-flag (Potentially re-transmitted message flag) is set. Failure of a peer is typically detected by the watchdog messages, but in some cases it is possible that a peer is not responding although watchdog messages are received. This could happen, for instance, when there is a Diameter relay agent or Diameter proxy agent between the system originating the Diameter messages and the final destination.
- The watchdog timer controls the frequency at which device-watchdog-request messages are transmitted to the peer, and is called the Tw timer in RFC 3539, *Authentication, Authorization and Accounting (AAA) Transport Profile*. A small timer results in a faster detection of a peer failure at the expense of generating more messages. The default is 30 seconds.

These timers can be configured at two levels: at Diameter peer policy level and at peer level. If configured at the peer level, then this value is taken for the specific peer, otherwise the configuration is taken from the Diameter peer policy level. If it is also not configured at the policy level, then the default values are used, which is 30 seconds for these three timers.

In case multiple peers are configured in the profile, a preference can be assigned to each of the peers. A lower preference value indicates a more preferred peer. Up to five peers can be configured, and all can be in the **I-Open** state, but the Diameter application will only select a single primary peer on a per application session basis. By default the preference is 50, and the selection on which peer is active and which one standby is shown in following CLI command:

```
*A:BNG-1# show aaa diameter-peer-policy "DSC.Geo.Red"
===============================================================================
Diameter Peer Policy : DSC.Geo.Red
===============================================================================
Last Mgmt Change      : 04/23/2014 15:18:24
-------------------------------------------------------------------------------
Diameter Base Values (config)
-------------------------------------------------------------------------------
Origin Host           : wlangw-2.DSC.Geo.Red.SRrealm
Origin Realm          : SRrealm
Connection Timer      : 30 (default)       Source Address     : 10.23.0.130
Transaction Timer     : 30 (default)       Router             : 10000
Watchdog Timer        : 30 (default)
Vendor Support        : 3GPP (default)
-------------------------------------------------------------------------------
Peer Name            Oper  PSM State     Susp  Cooldown  Pref  Order  Pri/Sec
-------------------------------------------------------------------------------
DSC.Simul            Yes   I-Open        No    -         10    1      Primary
DSC.26.206           Yes   I-Open        No    -         20    2      Secondary
===============================================================================
```

The configuration corresponding to the above show command is as follows.

```
*A:BNG-1# configure aaa diameter-peer-policy "DSC.Geo.Red"
*A:BNG-1>config>aaa>diam-peer-plcy# info
----------------------------------------------
            applications gx
            origin-host "wlangw-2.DSC.Geo.Red.SRrealm"
            origin-realm "SRrealm"
            router 10000
            source-address 10.23.0.130
            peer "DSC.Simul" create
                address 10.55.2.2
                destination-realm "simul.org"
                preference 10
                no shutdown
            exit
            peer "DSC.26.206" create
                address 10.40.11.2
                destination-realm "Tc3eRealm"
                preference 20
                no shutdown
            exit
----------------------------------------------
```

Which redundancy features are supported as well as redundancy behavior is Diameter application specific.

The transport configuration is part of the Diameter peer policy and is configured per peer. SR OS uses TCP as transport and the TCP destination port number is configurable. By default the standard port 3868 is used.

```
*A:BNG-1# configure aaa diameter-peer-
policy "DSC.26.206" peer "DSC.26.206" transport
  - transport tcp port <port>
  - no transport

 <tcp>                : keyword
 <port>               : [1..65535]
```

The source port is randomly chosen from the ephemeral port-range.

Troubleshooting

Statistics of each peer can be displayed as follows:

```
*A:BNG-1# show aaa diameter-peer-policy "DSC.26.206" peer "DSC.26.206" statistics
===============================================================================
Diameter Peer Policy : DSC.26.206 (statistics)
===============================================================================
Diameter Peer             : DSC.26.206
time statistics cleared   : 04/18/2014 07:52:28
```

```
-------------------------------------------------------------------------------
Client initiated tx/rx                  Server initiated tx/rx
-------------------------------------------------------------------------------
TCP Send Failed           : 0           TCP Send Failed           : 0
Diam Rx Drop Count (Resps): 0           Diam Rx Drop Count (Reqs) : 0
Diam Tx Requests          : 3           Diam Rx Requests          : 94
Diam Rx Responses         : 3           Diam Tx Responses         : 94
Pending Messages          : 0
Request Timeouts          : 0
-------------------------------------------------------------------------------
Diameter message breakdown
-------------------------------------------------------------------------------
CCR initial Tx            : 1           CCA initial Rx            : 1
CCR update Tx             : 1           CCA update Rx             : 1
CCR terminate Tx          : 1           CCA terminate Rx          : 1
CER Tx                    : 0           CEA Rx                    : 0
DWR Tx                    : 0           DWA Rx                    : 0
DWR Rx                    : 94          DWA Tx                    : 94
ASR Rx                    : 0           ASA Tx                    : 0
RAR Rx                    : 0           RAA Tx                    : 0
DPR Tx                    : 0           DPA Rx                    : 0
DPR Rx                    : 0           DPA Tx                    : 0
===============================================================================
```

The above command shows several statistics including the number of transmitted and received messages per message type. There is a command to clear the above counters, for instance:

```
clear aaa diameter-peer-policy "DSC.26.206" peer "DSC.26.206" statistics
```

Furthermore, debug at peer level of all Diameter message types are available:

```
debug
    diameter
        detail-level high
        no dest-realm
        diameter-peer DSC.26.206 psm-events
        no diameter-policy
        message-type ccr cca cer cea dwr dwa dpr dpa rar raa asr asa
        no origin-realm
    exit
exit
```

The above debug example will display the detailed output of all Diameter messages sent to and received from peer **DSC.26.206**. A shorter command to obtain the same output is the **debug diameter message-type all** command. Some of the Diameter message types are application specific with certain overlap between the applications. In other words, not all Diameter message types are used by all applications. The message types are:

- ccr credit-control-request
- cca credit-control-answer

- cer capability-exchange-request
- cea capability-exchange-answer
- dwr watchdog-request
- dwa watchdog-answer
- dpr disconnect-peer-request
- dpa disconnect-peer-answer
- rar re-auth-request
- raa re-auth-answer
- asr abort-session-request
- asa abort-session-answer

Note that debug of a request message and the corresponding answer message requires enabling debug for 2 message–types. Common practice is to enable debug for all message types, or for all message types except for the watchdog messages because typically these messages do not contain much interesting information. Including the periodic DWR and DWA messages would make the debug output harder to read.

Every Diameter application supports at least the CER/CEA messages. An example output of CER/CEA messages is shown below. In the CER sent by SR OS, the attributes origin-host, origin-realm, host-ip-addr, and auth-appl-id are coming from the diameter policy configuration. Note that a CER has no destination-host. Other information in the CER is the product-name (set to **SR-OS**) and firmware-revision (set to **1203** in the example below indicating that this debug trace is taken from 12.0R3). The CEA is received from the PCRF, and contains similar information.

```
5 2014/06/26 14:27:32.79 CET MINOR: DEBUG #2001 vprn10000 DIAMETER
"DIAMETER: Message Transmission
CER from [DSC.26.206, DSC.26.206] to 10.40.11.2:3868
  Header
    ver 1  len 196  flags R-------  code 257
    app-id 0  hbh-id 8652  e2e-id 659536111
  AVPs
    origin-host (264) -M------ [24]
      data [16] (DiameterIdentity) : wlangw-2.SRrealm
    origin-realm (296) -M------ [15]
      data [7] (DiameterIdentity) : SRrealm
    host-ip-addr (257) -M------ [14]
      data [6] (Address) : ipv4 10.23.0.130
    vendor-id (266) -M------ [12]
      data [4] (Unsigned32) : 6527
    product-name (269) -------- [13]
      data [5] (UTF8String) : SR-OS
    auth-appl-id (258) -M------ [12]
      data [4] (Unsigned32) : 16777238 : Gx
    supported-vendor-id (265) -M------ [12]
      data [4] (Unsigned32) : 6527
    supported-vendor-id (265) -M------ [12]
```

```
      data [4] (Unsigned32) : 10415
    supported-vendor-id (265) -M------ [12]
      data [4] (Unsigned32) : 13019
    vend-specific-appl-id (260) -M------ [32]
      data [24] (Grouped)
        vendor-id (266) -M------ [12]
          data [4] (Unsigned32) : 10415
        auth-appl-id (258) -M------ [12]
          data [4] (Unsigned32) : 16777238 : Gx
    firmware-revision (267) -------- [12]
      data [4] (Unsigned32) : 1203

  01 00 00 c4  80 00 01 01  00 00 00 00  00 00 21 cc
  27 4f b8 ef  00 00 01 08  40 00 00 18  77 6c 61 6e
  67 77 2d 32  2e 53 52 72  65 61 6c 6d  00 00 01 28
  40 00 00 0f  53 52 72 65  61 6c 6d 00  00 00 01 01
  40 00 00 0e  00 01 0a 17  00 82 00 00  00 00 01 0a
  40 00 00 0c  00 00 19 7f  00 00 01 0d  00 00 00 0d
  53 52 2d 4f  53 00 00 00  00 00 01 02  40 00 00 0c
  01 00 00 16  00 00 01 09  40 00 00 0c  00 00 19 7f
  00 00 01 09  40 00 00 0c  00 00 28 af  00 00 01 09
  40 00 00 0c  00 00 32 db  00 00 01 04  40 00 00 20
  00 00 01 0a  40 00 00 0c  00 00 28 af  00 00 01 02
  40 00 00 0c  01 00 00 16  00 00 01 0b  00 00 00 0c
  00 00 04 b3
"

6 2014/06/26 14:27:32.79 CET MINOR: DEBUG #2001 vprn10000 DIAMETER
"DIAMETER: Message Reception
CEA from 10.40.11.2:3868 to [DSC.26.206, DSC.26.206]
  Header
    ver 1  len 776  flags --------  code 257
    app-id 0  hbh-id 8652  e2e-id 659536111
  AVPs
    origin-host (264) -M------ [25]
      data [17] (DiameterIdentity) : stefaan.Tc3eRealm
    origin-realm (296) -M------ [17]
      data [9] (DiameterIdentity) : Tc3eRealm
    result-code (268) -M------ [12]
      data [4] (Unsigned32) : 2001 : DIAM_RESCODE_SUCCESS
    host-ip-addr (257) -M------ [14]
      data [6] (Address) : ipv4 10.40.11.2
    vendor-id (266) -M------ [12]
      data [4] (Unsigned32) : 637
    product-name (269) -------- [36]
      data [28] (UTF8String) : Nokia 5780 DSC (PS)
    origin-state-id (278) -M------ [12]
      data [4] (Unsigned32) : 1364308599
    firmware-revision (267) -------- [12]
      data [4] (Unsigned32) : 600450000
    auth-appl-id (258) -M------ [12]
      data [4] (Unsigned32) : 16777217 :
    auth-appl-id (258) -M------ [12]
      data [4] (Unsigned32) : 16777266 :
    auth-appl-id (258) -M------ [12]
      data [4] (Unsigned32) : 1 :
    auth-appl-id (258) -M------ [12]
      data [4] (Unsigned32) : 16777267 :
    auth-appl-id (258) -M------ [12]
```

```
    data [4] (Unsigned32) : 16777302 :
  auth-appl-id (258) -M------ [12]
    data [4] (Unsigned32) : 16777303 :
  auth-appl-id (258) -M------ [12]
    data [4] (Unsigned32) : 16777236 :
  auth-appl-id (258) -M------ [12]
    data [4] (Unsigned32) : 16777238 : Gx
  auth-appl-id (258) -M------ [12]
    data [4] (Unsigned32) : 111 :
  supported-vendor-id (265) -M------ [12]
    data [4] (Unsigned32) : 28458
  supported-vendor-id (265) -M------ [12]
    data [4] (Unsigned32) : 10415
  supported-vendor-id (265) -M------ [12]
    data [4] (Unsigned32) : 12951
  supported-vendor-id (265) -M------ [12]
    data [4] (Unsigned32) : 9
  supported-vendor-id (265) -M------ [12]
    data [4] (Unsigned32) : 7898
  supported-vendor-id (265) -M------ [12]
    data [4] (Unsigned32) : 5535
  supported-vendor-id (265) -M------ [12]
    data [4] (Unsigned32) : 637
  vend-specific-appl-id (260) -M------ [32]
    data [24] (Grouped)
      vendor-id (266) -M------ [12]
        data [4] (Unsigned32) : 10415
      auth-appl-id (258) -M------ [12]
        data [4] (Unsigned32) : 16777236 :
  vend-specific-appl-id (260) -M------ [32]
    data [24] (Grouped)
      vendor-id (266) -M------ [12]
        data [4] (Unsigned32) : 10415
      auth-appl-id (258) -M------ [12]
        data [4] (Unsigned32) : 16777267 :
  vend-specific-appl-id (260) -M------ [32]
    data [24] (Grouped)
      vendor-id (266) -M------ [12]
        data [4] (Unsigned32) : 10415
      auth-appl-id (258) -M------ [12]
        data [4] (Unsigned32) : 16777238 : Gx
  vend-specific-appl-id (260) -M------ [32]
    data [24] (Grouped)
      vendor-id (266) -M------ [12]
        data [4] (Unsigned32) : 9
      auth-appl-id (258) -M------ [12]
        data [4] (Unsigned32) : 16777238 : Gx
  vend-specific-appl-id (260) -M------ [32]
"
```

It is also possible to debug all messages from a specific Diameter policy, origin realm, or destination realm.

Note that these debug commands only show Diameter messages but no TCP messages like TCP-SYN. TCP layer issues must be debugged differently. For instance when there is a routing issue between client and server, then typically the state of the Diameter peer is **Wait-Conn-Ack**:

```
*A:BNG-1# show aaa diameter-peer-policy "DSC.26.206"
===============================================================================
Diameter Peer Policy : DSC.26.206
===============================================================================
Last Mgmt Change      : 06/16/2014 13:21:26
-------------------------------------------------------------------------------
Diameter Base Values (config)
-------------------------------------------------------------------------------
Origin Host           : wlangw-2.SRrealm
Origin Realm          : SRrealm
Connection Timer      : 30 (default)       Source Address     : 10.23.0.130
Transaction Timer     : 30 (default)       Router             : 10000
Watchdog Timer        : 10
Vendor Support        : 3GPP (default)
-------------------------------------------------------------------------------
Peer Name            Oper  PSM State     Susp  Cooldown  Pref  Order  Pri/Sec
-------------------------------------------------------------------------------
DSC.26.206           Yes   Wait-Conn-Ack No    Pending   50    -      -
===============================================================================
```

Wait-Conn-Ack means that the client has sent a TCP-SYN but no SYN-ACK is coming back. If the state is **Closed**, the client is no longer listening for a SYN-ACK and a new attempt to bring up the transport connection is made when the connection timer expires:

```
*A:BNG-1# show aaa diameter-peer-policy "DSC.26.206"
===============================================================================
Diameter Peer Policy : DSC.26.206
===============================================================================
Last Mgmt Change      : 06/16/2014 13:21:26
-------------------------------------------------------------------------------
Diameter Base Values (config)
-------------------------------------------------------------------------------
Origin Host           : wlangw-2.SRrealm
Origin Realm          : SRrealm
Connection Timer      : 30 (default)       Source Address     : 10.23.0.130
Transaction Timer     : 30 (default)       Router             : 10000
Watchdog Timer        : 10
Vendor Support        : 3GPP (default)
-------------------------------------------------------------------------------
Peer Name            Oper  PSM State     Susp  Cooldown  Pref  Order  Pri/Sec
-------------------------------------------------------------------------------
DSC.26.206           Yes   Closed        No    Pending   50    -      -
===============================================================================
```

The countdown of the connection timer can be seen with this command:

```
*A:BNG-1# show aaa diameter-peer-policy "DSC.26.206" peer "DSC.26.206"
===============================================================================
Diameter Peer Policy : DSC.26.206
===============================================================================
Diameter Peer          : DSC.26.206
Peer IP address        : 10.40.11.2
Last Mgmt Change       : 06/26/2014 14:27:31
-------------------------------------------------------------------------------
Peer Runtime Values (main)
```

```
-------------------------------------------------------------------------------
Peer Table Entry       : DSC.26.206::DSC.26.206
Peer Operational       : Yes              Watchdog Algorithm Active   : No
Peer State Machine     : Closed           Watchdog Answer Pending     : No
Connection Timer  (Tc) : 16               Connection Suspended        : No
Transaction Timer (Tt) : -                Cooldown Sequence Pending   : Yes
Watchdog Timer    (Tw) : -                Cooldown Sequence Active    : No
Primary/Secondary Peer : -                Cooldown Sequence Progress  : -
                                          Peer Removal Pending        : No
===============================================================================
```

In the above example, a new attempt to bring up the TCP session will be made in 16 seconds.

Events

Three events are defined for Diameter:

```
*A:BNG-1# show log event-control "diameter"
===============================================================================
Log Events
===============================================================================
Application
 ID#    Event Name                          P   g/s     Logged     Dropped
-------------------------------------------------------------------------------
   2001 tmnxDiamPolicyPeerStateChange    MI  gen       8095           0
   2002 tmnxDiamAppMessageDropped        MI  gen          6           0
   2003 tmnxDiamAppSessionFailure        MI  gen          0           0
===============================================================================
```

Trap **tmnxDiamPolicyPeerStateChange** is generated for all changes in the state of the Diameter peer. The second trap, **tmnxDiamAppMessageDropped** is generated when the system drops a Diameter message because it is malformed. Failures in the Diameter application sessions are reported in the trap **tmnxDiamAppSessionFailure**. Note that each Diameter application can have its own specific behavior for each of these traps. These traps are generated when a log is created from security:

```
configure log log-id 88
       from security
       to …
```

Conclusion

Diameter is an alternative to RADIUS. Although it is mainly used by mobile operators, it is finding its way in fixed access networks. Diameter peering provides reliable and secure transport with peer redundancy. Its functionality is defined in a base Diameter protocol specified in RFC3588. Various applications can be layered on top of base Diameter and they can utilize the robust transport capabilities that Diameter provides.

Flexible Authentication Model in ESM

This chapter provides information about Flexible Authentication Models in ESM.

Topics in this chapter include:

Applicability

This example is applicable to Routed Central Office (RCO) model on 7750 SR-7/12/12e, 7750 SR-c4/12 and 7450 ESS 7/12 in mixed-mode with IOM3-XP or IMM.

The configuration was tested on release 11.0R2 in a single-homed scenario.

Overview

The flexible authentication model for IPoE and PPPoE subscribers allows for mixing of configuration parameters obtained during the authentication phase from different sources: LUDB, RADIUS or Local User Database (LUDB), RADIUS or DHCP options that can be populated via a custom Python script. In case the same parameter is available from multiple sources, a priority mechanism is enforced whereby the parameter received from a higher priority source overrides the parameters received from the lower priority source in the following priority LUDB to RADIUS to Python.

In this example we will configure a dual-stack IPoE and a dual stack PPPoE host using 4 different methods to obtain their configuration parameters. The setup will utilize a single 7x50 BNG node with a locally configured DHCP server and LUDB as well as an external RADIUS server. Subscriber hosts are instantiated on managed (dynamic) SAPs.

The subscriber configuration parameters are in general divided into two categories:

- IP addressing parameters of the host — IPv4/v6 address/prefix, DNS servers, IPv4 default-gateway, IPv4 subnet-mask, IPv4/v6 address pool name, DHCPv4/v6 lease times, etc.

- Non IP addressing parameters of the host — Subscriber hosts strings are used to associate the subscriber-host with the desired level of service (sub/sla-profiles, inter-dest-id string, etc); managed routes are used for routing purposes to/from the host; etc.

The following four scenarios will be examined:

1. DHCP relay case (IP address is assigned via local DHCP server) with NO authentication. See DHCP Relay Case with No Authentication.
2. DHCP relay case (IP address is assigned via local DHCP server) with LUDB + RADIUS authentication. See DHCP Relay Case with LUDB + RADIUS Authentication.

   RADIUS provides: sub/sla-profile strings and a framed IPv4 route.

   LUDB provides: IP address pool, inter-dest-id string for Vport assignment, msap-defaults (routing context parameters and msap-policy).
3. IP proxy case (IP address is assigned via RADIUS) with LUDB + RADIUS authentication. IP Proxy Case with LUDB + RADIUS Authentication

   RADIUS provides: IP addresses and related parameters (DNS server, IPv4 default-gateway, etc), inter-dest-id string for Vport assignment and a framed route.

Note: IPv6 lease-times are provided under the group-interface.

   LUDB provides: sub/sla-profile strings and msap-defaults (routing context parameters and msap-policy).
4. IP proxy case (IP address is assigned via LUDB) with LUDB + RADIUS authentication. IP Proxy Case with LUDB + RADIUS Authentication

   RADIUS provides: sub/sla-profile strings and a framed IPv4 route.

   LUDB provides IP addresses and related parameters (DNS server, IPv4 default-gateway, etc), inter-dest-id string for Vport assignment and msap-defaults (routing context parameters and msap-policy).

In cases 2-4, the domain-name *alu-domain* is appended to the IPoE and PPPoE username in LUDB, just before RADIUS authentication takes place.

Configuration

The topology is shown in Figure 149.

Figure 149 **Topology**

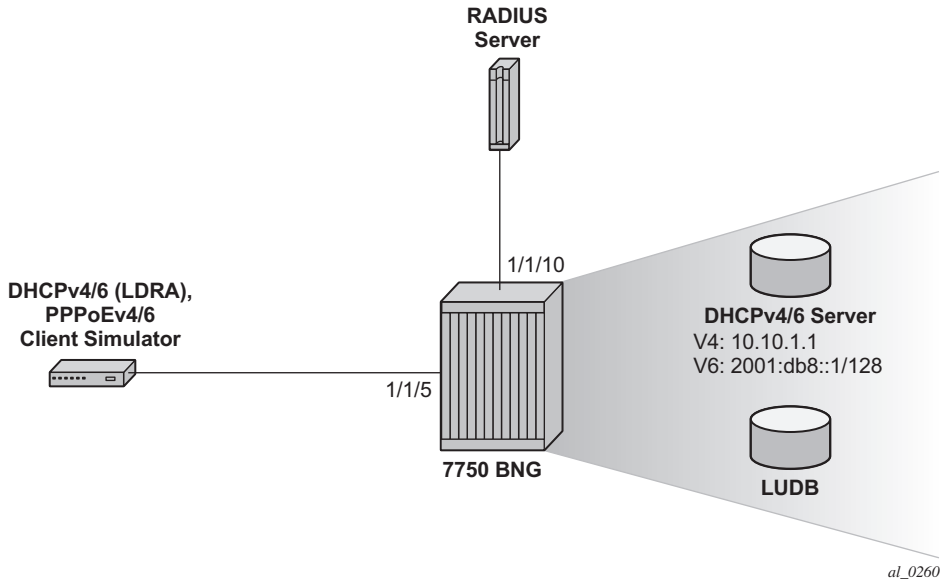

There is a common part of the configuration that applies uniformly across all four examined scenarios. This common part is outlined below and will not be repeated again when we describe more specific cases. It is assumed that the more specific cases also contain this common part of the configuration.

Common Configuration Examples

Access Ethernet Port with QinQ Encapsulation

The following output displays a configuration example.

```
configure port 1/1/5
     ethernet
        mode access
        encap-type qinq
     exit
     no shutdown
```

Capture SAP

A capture SAP is used to dynamically detect the VLAN id(s) in incoming DHCP/PPPoE packets (triggering packets) and conditionally instantiate the managed (dynamic) SAP. LUDB must be configured under the capture SAP to authorize the user accessing the capture SAP. The LUDB may contain additional parameters needed to setup the subscriber, it can point the subscriber to the RADIUS server for additional parameters or it may contain a default subscriber-host entry without any configuration parameters.

In this case the **msap-defaults** under the capture SAP is used to select the routing context where the msap is created. **msap-defaults** can be also configured in the LUDB or be supplied via RADIUS.

PPPoE policy and *msap policy* are used to define PPPoE and SAP level parameters. Since the (dynamic) SAP does not exist at the time when the initial DHCP/PPPoE packets are received, the PPPoE/SAP level parameters are taken from the PPPoE/msap policy under the capture SAP. For example, those parameters are used in the PPP PADx/LCP/Authentication setup phase, they define default subscriber host strings, maximum number of subscriber hosts per SAP, the anti-spoofing mode, etc.

```
configure service vpls 2
    sap 1/1/5:17.* capture-sap create
     description "open DHCP model testing"
     trigger-packet dhcp dhcp6 pppoe
     dhcp-user-db "open-dhcp"
     dhcp6-user-db "open-dhcp"
     pppoe-policy "pppoe_pol"
     pppoe-user-db "open-dhcp"
     msap-defaults
         group-interface "open-auth"
     policy "msaps"
     service 10
     exit
exit
```

auto-sub-id

The **auto-sub-id-key** command can be used in situations where the more specific *subscriber-id* string is not returned from LUDB or RADIUS. In this case, the auto subscriber-id for IPoE hosts is set to the circuit-id while for PPPoE hosts the auto subscriber-id is set to the circuit-id plus session-id separated by the “|”delimiter which is inserted by default.

```
configure subscriber-mgmt auto-sub-id-key
    ipoe-sub-id-key circuit-id
    ppp-sub-id-key circuit-id session-id
```

PPPoE Policy

There is a maximum of PPPoE sessions per MAC on a managed SAP. The default is 1 but is increased here to 10.

```
configure subscriber-mgmt ppp-policy "pppoe_pol"
     ppp-mtu 1400
     max-sessions-per-mac 10
```

MSAP Policy

The MSAP policy defines the anti-spoofing mode which is in this particular example set to next-hop MAC (nh-mac). It also defines the default subscriber management parameters in case they are not supplied via LUDB or RADIUS.

```
configure subscriber-mgmt msap-policy <msap-policy-name> create
    sub-sla-mgmt
        def-sub-id use-auto-id
        def-sub-profile "default-sub-profile"
        def-sla-profile "default-sla"
        sub-ident-policy "sub_ident_pol"
        multi-sub-sap limit 500
    exit
    ies-vprn-only-sap-parameters
        anti-spoof nh-mac
    exit
```

subscriber-interface Configuration

The following output displays a subscriber interface configuration.

```
configure service vprn 10
    subscriber-interface "sub-int-1" create
        allow-unmatching-subnets          Support for 'un-numbered'(1) IPv4 clients.
        address 10.12.0.1/24            Default gateway for IPv4 'numbered' clients.
        ipv6
            delegated-prefix-len 56      Fixed delegated prefix length for IA-PD.
            allow-unmatching-prefixes     Support for 'un-numbered' IPv6 clients.
        exit
        group-interface "open-auth" create
            ipv6
                router-advertisements
                    managed-configuration       Hint to the client to use DHCPv6.
                    no shutdown                Enabling Router-Advertisements.
                exit
                dhcp6
                    user-db "open-dhcp"   Must be the same as under the capture-sap.
```

```
                exit
            exit
            arp-populate      ARP table is populated based on the lease-state table.
            dhcp
                trusted           Accept DHCP packets on this group-interface.
                lease-
populate 1000 Max number of DHCPv4 clients on each SAP of this grp-intf.
                user-db "open-dhcp"     Must be the same as under the capture-sap.
            exit
            pppoe
                policy "pppoe_pol"
                session-limit 1000
                sap-session-limit 1000
                user-db "open-dhcp"      Must be the same as under the capture-sap.
                no shutdown
            exit
```

Note:

1. numbered/unnumbered subscriber-hosts. Refer to the DHCP/PPPoE clients whose assigned IP address is outside of any IP subnet/prefix configured under the subscriber-interface.

Specific Configuration Parts

DHCP Relay Case with No Authentication

The IP address is assigned via local DHCP server. The LUDB is accessed even in the scenario without authentication. There must be a default host LUDB entry present that will match on any value specified in the match-list criteria. The LUDB is accessed from the capture SAP (part of the common configuration).

```
    configure subscriber-mgmt local-user-db "open-dhcp" create
        dhcp
            match-list circuit-id  Host matching based on circuit-
id in DHCP packets.
            host "default" create
                no shutdown
        exit
        exit
        ppp
            match-list username              Host matching based on PPPoE username.
            host "default" create    force-ipv6cp      Explicitly enabled IPCPv6.
                no shutdown
            exit
        exit
        no shutdown
```

Once the routing context (service id and group-interface) is determined as defined under the capture SAP defaults (part of the common configuration), the DHCP/ PPPoE requests are served according to the group-interface configuration. The IP address request is relayed to the DHCPv4/v6 server. Since the LUDB does not provide a pool name, the **gi-address** and the **link-address** is used by the DHCP relay/server to select the pool from which the IP address will be assigned.

```
    configure service vprn 10 subscriber-interface "sub-int-1" group-
interface "open-auth" create
        ipv6
            dhcp6
            relay                                   DHCPv6 relay configuration.
                link-address 2001:DB8:1::
                server 2001:DB8::1                  DHCPv6 server IPv6 address.
                client-applications dhcp ppp
                no shutdown
            exit
        exit
    exit
    dhcp
        server 10.10.1.1                                    DHCPv4 server IP address.
        client-applications dhcp ppp
        gi-address 10.12.0.1
        no shutdown
    exit
```

DHCPv4/v6 servers are locally configured in the 7x50 and attached to a loopback interface.

```
    configure service vprn 10 interface "loop-dhcp-srvr"
        address 10.10.1.1/24  IPv4       Address to which is DHCPv4 server attached.
        ipv6
            address 2001:DB8::1/
128   IPv6 address to which is DHCPv6 server attached.
            local-dhcp-server "v6"     Attaching DHCPv6 server to the loopback intf.
            exit
        local-dhcp-
server "local"       Attaching DHCPv4 server to the loopback intf.
        loopback
```

In the local DHCP servers two pools are defined:

- LUDB — To be used for IP address assignment when LUDB returns the pool name.
- Gi-addr — To be used when gi-address/link-address are used to select the pool for IP address assignment.

Lease times for IPv4 and IPv6 are configured in the local DHCP server which is used only in the relay case (when the IP address is supplied via DHCP server and not through RADIUS or the LUDB).

```
configure service  vprn 10
    dhcp
        local-dhcp-server "local" create
        use-gi-address                The gi-address can be used to select the pool.
          use-pool-from-client         The pool name can be explicitly provided.
          pool "ludb" create          The pool used when LUDB provides the pool name.
          options
                dns-server 172.16.16.16 172.16.16.17
               lease-time hrs 1                                DHCPv4 lease time.
             exit
            subnet 10.10.0.0/24 create
               options
                 subnet-mask 255.255.255.0
                   default-router 10.10.0.1
                 exit
                address-range 10.10.0.100 10.10.0.200
             exit
          exit
           pool "gi-addr" create               Pool selected based on the gi-address.
               options
                 dns-server 172.16.16.16 172.16.16.17
                 lease-time hrs 1                                 DHCPv4 lease time.
                exit
                 subnet 10.12.0.0/24 create
                options
                 subnet-mask 255.255.255.0
                  default-router 10.12.0.1
                 exit
                  address-range 10.12.0.100 10.12.0.200
                 exit
             exit
          no shutdown
           exit
          exit
        dhcp6
           local-dhcp-server "v6" create
                use-link-address
                  use-pool-from-client
                  pool "ludb" create
                 prefix 2001:DB8:10::/48 pd wan-host create
                     preferred-lifetime min 30
                          rebind-timer min 20
                          renew-timer min 15
                          valid-lifetime hrs 1                    DHCPv6 lease time.
                          options
                           dns-server 2001:DB8::1000 2001:DB8::1001
                         exit
                      exit
                    exit

                    pool "gi-addr" create
                      prefix 2001:DB8:30::/48 pd wan-host create
                        preferred-lifetime min 30
                             rebind-timer min 20
                             renew-timer min 15
                             valid-lifetime hrs 1                 DHCPv6 lease time.
                             options
                              dns-server 2001:DB8::1000 2001:DB8::1001
                             exit
```

```
                    exit
                exit
            no shutdown
```

Default sub/sla-profiles, from the msap-policy, are used (part of the common configuration).

```
configure  subscriber-mgmt sla-profile "default-sla"
description "default SLA profile"
      host-limit 3

configure  subscriber-mgmt sub-profile "default-sub-profile"
     description "default SUB profile"
     egress
          agg-rate-limit 1000
      exit
```

Show Commands

The following command shows that the default sub/sla-profiles are in use, that the IP addresses are selected from the gi-addr pool in local DHCP server and that the subscriber-id is set to circuit-id for the IPoE subscriber-host and to *username*|*session-id* combination for the PPPoE subscriber-host.

```
A:BNG-1# show service  active-subscribers
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber open-dhcp (default-sub-profile)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[1/1/5:17.10] - sla:default-sla
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.12.0.101
                00:00:65:17:10:01 N/A       DHCP
2001:DB8:30::1/128
                00:00:65:17:10:01 N/A       IPoE-DHCP6
2001:DB8:30:100::/56
                00:00:65:17:10:01 N/A       IPoE-DHCP6
-------------------------------------------------------------------------------
Subscriber open-pppoe|2 (default-sub-profile)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[1/1/5:17.11] - sla:default-sla
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.12.0.100
```

```
               00:00:65:17:11:02 2         IPCP
2001:DB8:30:1::1/128
               00:00:65:17:11:02 2         PPP-DHCP6
2001:DB8:30:200::/56
               00:00:65:17:11:02 2         PPP-DHCP6
-------------------------------------------------------------------------------
Number of active subscribers : 2
```

The following command shows more details about the subscriber-host, such as the group-interface, address origin, acct-session-id, etc. Even though there are only two dual-stack hosts (one IPoE and one PPPoE), each of them has 3 IP addresses that show up as different hosts.

For the purpose of brevity, the output for only two IP hosts are shown, one with an IPv4 address and one with an IPv6 address. The remaining IP addresses/prefixes are not shown since the output follows the same logic.

```
A:BNG-1# show service id 10 subscriber-hosts detail
===============================================================
Subscriber Host table
===============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID Origin       Fwding State
---------------------------------------------------------------
[1/1/5:17.10]          open-dhcp
  10.12.0.101
    00:00:65:17:10:01    N/A       DHCP         Fwding
---------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : open-auth
Sub Profile           : default-sub-profile
SLA Profile           : default-sla
App Profile           : N/A
Egress Q-Group        : policer-output-queues
Egress Vport          : N/A
Acct-Session-Id       : D897FF0000000F51DBC5A7
Acct-Q-Inst-Session-Id: D897FF0000001051DBC5A7
Address Origin        : DHCP
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
---------------------------------------------------------------
 [1/1/5:17.11]          open-pppoe|2
  2001:DB8:30:1::1/128
    00:00:65:17:11:02    2         PPP-DHCP6    Fwding
---------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : open-auth
Sub Profile           : default-sub-profile
SLA Profile           : default-sla
App Profile           : N/A
Egress Q-Group        : policer-output-queues
Egress Vport          : N/A
Acct-Session-Id       : D897FF0000001351DBC5BA
Acct-Q-Inst-Session-Id: D897FF0000000E51DBC59C
```

```
Address Origin        : DHCP
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
-------------------------------------------------------------
Number of subscriber hosts : 6             The remaining 4 hosts are not shown here…
=============================================================
```

The following command shows that there are no sub/sla-profile strings assigned to the subscriber. Instead the default sub/sla-profiles from the msap-policy are used.

The IP address is assigned by the DHCP server which also supplied the def-gw information, DNS servers, the net-mask and the lease time.

The circuit-id and the subscriber-id are set to the same value.

```
A:BNG-1# show service id 10 dhcp lease-state detail
===============================================================================
DHCP lease states for service 10
===============================================================================
Service ID           : 10
IP Address           : 10.12.0.101
Client HW Address    : 00:00:65:17:10:01
Subscriber-interface : sub-int-1
Group-interface      : open-auth
SAP                  : [1/1/5:17.10]
Up Time              : 0d 00:04:01
Remaining Lease Time : 0d 00:56:00
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "open-dhcp"
Sub-Profile-String   : ""
SLA-Profile-String   : ""
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : ""
Category-Map-Name    : ""

Lease Info origin    : DHCP

Ip-Netmask           : 255.255.255.0
Broadcast-Ip-Addr    : N/A
Default-Router       : 10.12.0.1
Primary-Dns          : 172.16.16.16
Secondary-Dns        : 172.16.16.17
Primary-Nbns         : N/A
Secondary-Nbns       : N/A

ServerLeaseStart     : 07/09/2013 01:11:19
ServerLastRenew      : 07/09/2013 01:11:19
ServerLeaseEnd       : 07/09/2013 02:11:19
Session-Timeout      : N/A
Lease-Time           : 0d 01:00:00
DHCP Server Addr     : 10.10.1.1
```

```
Relay Agent Information
  Circuit Id          : open-dhcp
  Remote Id           : ipoe-v6
Radius User-Name      : ""
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
```

Then there is a similar command used for DHCPv6 lease-state details.

For the purpose of brevity, the output for only two IPv6 leases is shown. The remaining two IPv6 leases are not shown since the output follows the same logic.

```
A:BNG-1# show service id 10 dhcp6 lease-state detail
===============================================================================
DHCP lease states for service 10
===============================================================================
Service ID            : 10
IP Address            : 2001:DB8:30::1/128
Client HW Address     : 00:00:65:17:10:01
Subscriber-interface  : sub-int-1
Group-interface       : open-auth
SAP                   : [1/1/5:17.10]
Up Time               : 0d 00:44:50
Remaining Lease Time  : 0d 00:45:10
Remaining SessionTime: N/A
Persistence Key       : N/A

Sub-Ident             : "open-dhcp"
Sub-Profile-String    : ""
SLA-Profile-String    : ""
App-Profile-String    : ""
Lease ANCP-String     : ""
Lease Int Dest Id     : ""
Category-Map-Name     : ""
Dhcp6 ClientId (DUID): 00030001000065171001
Dhcp6 IAID            : 0
Dhcp6 IAID Type       : non-temporary
Dhcp6 Client Ip       : FE80::200:65FF:FE17:1001
Primary-Dns           : N/A
Secondary-Dns         : N/A
Pool Name             : ""
Dhcp6 Server Addr     : 2001:DB8::1
Dhcp6 ServerId (DUID): 00030001d897ff000000
Dhcp6 InterfaceId     : open-dhcp
Dhcp6 RemoteId        : 0000ipoe-v6

Lease Info origin     : DHCP

ServerLeaseStart      : 07/09/2013 01:11:28
ServerLastRenew       : 07/09/2013 01:41:28
ServerLeaseEnd        : 07/09/2013 02:41:28                    One hour lease time.
Session-Timeout       : N/A
Radius User-Name      : ""
-------------------------------------------------------------------------------
Service ID            : 10
IP Address            : 2001:DB8:30:1::1/128
```

```
Client HW Address    : 00:00:65:17:11:02
Subscriber-interface : sub-int-1
Group-interface      : open-auth
SAP                  : [1/1/5:17.11]
Up Time              : 0d 00:44:40
Remaining Lease Time : 0d 00:45:20
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "open-pppoe|2"
Sub-Profile-String   : ""
SLA-Profile-String   : ""
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : ""
Category-Map-Name    : ""
Dhcp6 ClientId (DUID): 00030001000065171102
Dhcp6 IAID           : 0
Dhcp6 IAID Type      : non-temporary
Dhcp6 Client Ip      : FE80::200:65FF:FE17:1102
Primary-Dns          : N/A
Secondary-Dns        : N/A
Pool Name            : ""
Dhcp6 Server Addr    : 2001:DB8::1
Dhcp6 ServerId (DUID): 00030001d897ff000000
Dhcp6 InterfaceId    : open-pppoe
Dhcp6 RemoteId       : N/A

Lease Info origin    : DHCP

ServerLeaseStart     : 07/09/2013 01:11:38
ServerLastRenew      : 07/09/2013 01:41:38
ServerLeaseEnd       : 07/09/2013 02:41:38
Session-Timeout      : N/A
Radius User-Name     : ""
-------------------------------------------------------------------------------
Number of lease states : 4               The remaining 2 leases are not shown here.
===============================================================================
```

DHCP Relay Case with LUDB + RADIUS Authentication

IP address is assigned via local DHCP server.

- RADIUS provides sub/sla-profile strings and a framed IPv4 route.
- LUDB provides IP address pool, inter-dest-id string for Vport assignment, msap-defaults (routing context parameters and msap-policy).

Vport aggregate rate limit and the port scheduler are now added to the physical port. The Vport is associated with the subscriber through the inter-dest-id string obtained via LUDB.

```
configure  port 1/1/5
     ethernet
```

```
        mode access
      encap-type qinq
       egress-scheduler-policy "port"
       access
            egress
               vport "open-dhcp" create
               agg-rate-limit 500
                   host-match dest "open-auth-vport" create
                exit
             exit
          exit
          no shutdown
        exit
```

The LUDB is used to assign the IP pool name (pool-name = ludb) and the inter-dest-id string (inter-dest-id = open-auth-vport) to the subscriber. The pool name is carried to the DHCP server via custom DHCP options [(82,9,13) in DHCPv4 and (17,1->wan_pool and 2->pfx_pool) in DHCPv6].

The domain name *alu-domain* is appended to the username (circuit-id = open-dhcp or username = open-pppoe) before an Access-Request message is sent to the RADIUS server which is configured in the authentication policy *open-dhcp*.

The inter-dest-id string taken from the LUDB is passed to the subscriber management module in the 7x50 via DHCP option 254 in DHCP ACK/Reply.

```
configure subscriber-mgmt local-user-db "open-dhcp"
local-user-db "open-dhcp" create
        dhcp
           match-list circuit-id
             host "open-dhcp" create
                host-identification
                   circuit-id string "open-dhcp"
                exit
                 address pool "ludb"
                 auth-policy "open-dhcp"
                auth-domain-name "alu-domain"
               identification-strings 254 create
                   inter-dest-id "open-auth-vport"
              exit
              msap-defaults
                   group-interface "open-auth"
                    policy "msaps"
                   service 10
                exit
                ipv6-wan-address-pool "ludb"
                ipv6-delegated-prefix-pool "ludb"
                no shutdown
             exit
          exit
           ppp
             match-list circuit-id username
           host "open-ppp" create
                 host-identification
                    username "open-pppoe"
```

```
                exit
                auth-policy "open-dhcp"
               address pool "ludb"
               password chap "ALU" hash2
              identification-strings 254 create
                    inter-dest-id "open-auth-vport"
               exit
              msap-defaults
                  group-interface "open-auth"
                  policy "msaps"
                   service 10
               exit
              ipv6-delegated-prefix-pool "ludb"
                ipv6-wan-address-pool "ludb"
              no shutdown
          exit
           exit
          no shutdown
exit
```

The *inter-dest-id* string taken from the LUDB is passed to the subscriber management module in the 7x50 via DHCPv4/v6 option 254 that is specified in the subscriber identification policy.

```
configure subscriber-mgmt sub-ident-policy "sub_ident_pol"
      strings-from-option 254
```

The RADIUS server is defined in the authentication policy. The domain name can be appended to the PPPoE subscriber host directly via the authentication-policy while for IPoE subscribers, the domain name is appended via the authentication-policy in conjunction with the LUDB. This can be verified in the output (shown later) of the **show service id 10 dhcp lease-state detail** and **show service id 10 dhcp6 lease-state detail** commands (on the “radius user-name” line).

```
configure subscriber-mgmt authentication-policy "open-dhcp"
password "ALU" hash2
      ppp-user-name append "alu-domain"
      radius-authentication-server
           server 1 address X.Y.Z.W secret "ALU" hash2
       exit
      user-name-format circuit-id append
       pppoe-access-method pap-chap
```

The RADIUS user configuration file uses the domain-name extension, as inserted by the 7x50 BNG, to authenticate the user:

```
open-dhcp@alu-domain   Cleartext-Password := "ALU"
            Alc-Subsc-Prof-Str = rad-sub,                Subscriber profile string.
            Alc-SLA-Prof-Str = rad-sla,                         SLA profile string.
            Framed-Route = "192.168.1.0/24 0.0.0.0"              Managed IPv4 route.
            Fall-Through = No
```

```
open-pppoe@alu-domain  Cleartext-Password := "ALU"
            Alc-Subsc-Prof-Str = rad-sub,
            Alc-SLA-Prof-Str = rad-sla,
            Framed-Route = "192.168.2.0/24 0.0.0.0",
            Fall-Through = No
```

DHCPv4/v6 servers are locally configured in the 7x50 and attached to a loopback interface:

```
configure  service  vprn  10 interface "loop-dhcp-srvr"
      address 10.10.1.1/24          IPv4 address to which is DHCPv4 server attached.
    ipv6
         address 2001:DB8::1/128    IPv6 address to which is DHCPv6 server attached.
         local-dhcp-server "v6"        Attaching DHCPv6 server to the loopback intf.
       exit
    local-dhcp-server "local"        Attaching DHCPv4 server to the loopback intf.
     loopback
```

Group-interface configuration. Note that common parts of the configuration as defined earlier, still apply:

```
configure service vprn 10 subscriber-interface "sub-int-1" group-interface "open-
auth"
      dhcp6
           user-db "open-dhcp"
              relay                                    DHCPv6 relay configuration.
              link-address 2001:DB8:30::
              server 2001:DB8::1
                client-applications dhcp ppp
                no shutdown
           exit
          exit
      dhcp                                             DHCPv4 relay configuration.
          option
            no circuit-id                  7x50 will not insert its own circuit-id.
          no remote-id                      7x50 will not insert its own remote-id.
           vendor-specific-option
                 pool-name
             exit
          exit
           server 10.10.1.1
        client-applications dhcp ppp
          no shutdown
       exit


exit
```

Lease times for IPv4 and IPv6 are configured in the local DHCP server. Lease times under the local DHCP server are used only in the relay case (when IP address is supplied via DHCP server and NOT RADIUS or LUDB). In the proxy case the lease times can be obtained via LUDB, RADIUS or group-interface.

```
configure service vprn 10
dhcp
           local-dhcp-server "local" create
              use-gi-address              gi-address can be used to select the pool.
           use-pool-from-client               pool name can be explicitly provided.
           pool "ludb" create             pool used when LUDB provides the pool name.
               options
             dns-server 172.16.16.16 172.16.16.17
                    lease-time hrs 1
                 exit
               subnet 10.10.0.0/24 create
                  options
                     subnet-mask 255.255.255.0
                       default-router 10.10.0.1
                     exit
                    address-range 10.10.0.100 10.10.0.200
                exit
            exit
             pool "gi-addr" create             pool selected based on the gi-address.
                options
                   dns-server 172.16.16.16 172.16.16.17
                   lease-time hrs 1
                  exit
              subnet 10.12.0.0/24 create
                    options
                      subnet-mask 255.255.255.0
                       default-router 10.12.0.1
                     exit
                      address-range 10.12.0.100 10.12.0.200
                     exit
                exit
               no shutdown
              exit
         exit
        dhcp6
             local-dhcp-server "v6" create
           use-link-address
           use-pool-from-client
           pool "ludb" create
               prefix 2001:DB8:10::/48 pd wan-host create
                   preferred-lifetime min 30
                   rebind-timer min 20
                   renew-timer min 15
                   valid-lifetime hrs 1
                   options
                       dns-server 2001:DB8::1000 2001:DB8::1001
                   exit
               exit
           exit
           pool "gi-addr" create
               prefix 2001:DB8:30::/48 pd wan-host create
                   preferred-lifetime min 30
                   rebind-timer min 20
                   renew-timer min 15
                   valid-lifetime hrs 1
                   options
                       dns-server 2001:DB8::1000 2001:DB8::1001
                   exit
               exit
```

```
            exit
            no shutdown
        exit
    exit
```

RADIUS sub/sla-profiles supplied via RADIUS are used:

```
configure subscriber-mgmt sla-profile "rad-sla"
     description "sla-profile obtained via RADIUS"
      host-limit 100
      egress
        qos 10 vport-scheduler
          exit
         ip-filter 10
exit
exit

configure subscriber-mgmt sub-profile "rad-sub"
      description "sub-profile obtained via RADIUS"
       egress
           agg-rate-limit 15000
       exit
exit
```

Show Commands

The following command shows that the rad-sub/sla-profiles, as supplied via RADIUS, are in use.

The IP addresses are selected from the pool-name LUDB in the local DHCP server. The subscriber-id is *circuit-id* for IPoE subscriber-host and the *username*|*session-id* combination for PPPoE subscriber host.

```
A:BNG-1#show service active-subscribers
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber open-dhcp (rad-sub)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[1/1/5:17.10] - sla:rad-sla
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.10.0.101
                                00:00:65:17:10:01 N/A       DHCP
2001:DB8:10:1::1/128
                00:00:65:17:10:01 N/A       IPoE-DHCP6
2001:DB8:10:200::/56
                00:00:65:17:10:01 N/A       IPoE-DHCP6
-------------------------------------------------------------------------------
```

```
Subscriber open-pppoe|3 (rad-sub)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[1/1/5:17.11] - sla:rad-sla
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.10.0.100
                00:00:65:17:11:02 3         IPCP
2001:DB8:10::1/128
                00:00:65:17:11:02 3         PPP-DHCP6
2001:DB8:10:100::/56
                00:00:65:17:11:02 3         PPP-DHCP6
-------------------------------------------------------------------------------
Number of active subscribers : 2
-------------------------------------------------------------------------------
```

The following command shows more details about the subscriber-host, such as the group-interface, vport, address origin, acct-session-id, etc. Vport is selected based on the *inter-dest-id* string supplied via the LUDB.

For the purpose of brevity, the output for only two IP addresses **hosts** is shown, one with an IPv4 address and one with an IPv6 address. The remaining IP addresses/ prefixes are not shown since the output follows the same logic.

```
A:BNG-1# show service id 10 subscriber-hosts detail
===============================================================
Subscriber Host table
===============================================================
Sap                     Subscriber
  IP Address
    MAC Address           PPPoE-SID Origin       Fwding State
---------------------------------------------------------------
[1/1/5:17.10]           open-dhcp
  10.10.0.101
    00:00:65:17:10:01    N/A       DHCP          Fwding
---------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : open-auth
Sub Profile           : rad-sub
SLA Profile           : rad-sla
App Profile           : N/A
Egress Q-Group        : policer-output-queues
Egress Vport          : open-dhcp
Acct-Session-Id       : D897FF0000000551D308B2
Acct-Q-Inst-Session-Id: D897FF0000000651D308B2
Address Origin        : DHCP
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
---------------------------------------------------------------
[1/1/5:17.11]           open-pppoe|3
  2001:DB8:10::1/128
    00:00:65:17:11:02    3         PPP-DHCP6     Fwding
---------------------------------------------------------------
Subscriber-interface  : sub-int-1
```

```
Group-interface       : open-auth
Sub Profile           : rad-sub
SLA Profile           : rad-sla
App Profile           : N/A
Egress Q-Group        : policer-output-queues
Egress Vport          : open-dhcp
Acct-Session-Id       : D897FF0000000351D308AF
Acct-Q-Inst-Session-Id: D897FF0000000251D308A9
Address Origin        : DHCP
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
-------------------------------------------------------------
```

The following command shows that the subscriber identity is set to *circuit-id* (plus session-id) as instructed by **auto-sub-id-key** command (subscriber-id string is not returned via the LUDB or RADIUS). The lease times are set to 1h as defined in the DHCP server. The username passed to RADIUS is a *circuit-id* or a *username* appended by the *alu-dmain* domain name.

```
A:BNG-1# show service id 10 dhcp lease-state  detail
===============================================================================
DHCP lease states for service 10
===============================================================================
Service ID           : 10
IP Address           : 10.10.0.101
Client HW Address    : 00:00:65:17:10:01
Subscriber-interface : sub-int-1
Group-interface      : open-auth
SAP                  : [1/1/5:17.10]
Up Time              : 0d 00:12:45
Remaining Lease Time : 0d 00:47:16
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "open-dhcp"
Sub-Profile-String   : "rad-sub"
SLA-Profile-String   : "rad-sla"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : "open-auth-vport"
Category-Map-Name    : ""

Lease Info origin    : DHCP

Ip-Netmask           : 255.255.255.0
Broadcast-Ip-Addr    : N/A
Default-Router       : 10.10.0.1
Primary-Dns          : 172.16.16.16
Secondary-Dns        : 172.16.16.17
Primary-Nbns         : N/A
Secondary-Nbns       : N/A

ServerLeaseStart     : 07/02/2013 10:06:58
ServerLastRenew      : 07/02/2013 10:06:58
ServerLeaseEnd       : 07/02/2013 11:06:58
Session-Timeout      : N/A
```

```
Lease-Time           : 0d 01:00:00
DHCP Server Addr     : 10.10.1.1

Relay Agent Information
  Circuit Id         : open-dhcp
  Remote Id          : ipoe-v6
Radius User-Name     : "open-dhcp@alu-domain"

Managed Routes       : 192.168.1.0/24           installed
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
```

For the purpose of brevity the output for only two IPv6 leases is shown. The remaining two IPv6 leases are not shown since the output follows the same logic.

```
A:BNG-1# show service id 10 dhcp6 lease-state  detail
===============================================================================
DHCP lease states for service 10
===============================================================================
Service ID           : 10
IP Address           : 2001:DB8:10::1/128
Client HW Address    : 00:00:65:17:11:02
Subscriber-interface : sub-int-1
Group-interface      : open-auth
SAP                  : [1/1/5:17.11]
Up Time              : 0d 00:13:00
Remaining Lease Time : 0d 00:47:00
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "open-pppoe|3"
Sub-Profile-String   : "rad-sub"
SLA-Profile-String   : "rad-sla"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : "open-auth-vport"
Category-Map-Name    : ""
Dhcp6 ClientId (DUID): 00030001000065171102
Dhcp6 IAID           : 0
Dhcp6 IAID Type      : non-temporary
Dhcp6 Client Ip      : FE80::200:65FF:FE17:1102
Primary-Dns          : N/A
Secondary-Dns        : N/A
Pool Name            : "ludb"
Dhcp6 Server Addr    : 2001:DB8::1
Dhcp6 ServerId (DUID): 00030001d897ff000000
Dhcp6 InterfaceId    : open-pppoe
Dhcp6 RemoteId       : N/A

Lease Info origin    : DHCP

ServerLeaseStart     : 07/02/2013 10:06:55
ServerLastRenew      : 07/02/2013 10:06:55
ServerLeaseEnd       : 07/02/2013 11:06:55
Session-Timeout      : N/A
Radius User-Name     : "open-pppoe@alu-domain"
```

```
-------------------------------------------------------------------------------
Service ID           : 10
IP Address           : 2001:DB8:10:1::1/128
Client HW Address    : 00:00:65:17:10:01
Subscriber-interface : sub-int-1
Group-interface      : open-auth
SAP                  : [1/1/5:17.10]
Up Time              : 0d 00:12:52
Remaining Lease Time : 0d 00:47:08
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "open-dhcp"
Sub-Profile-String   : "rad-sub"
SLA-Profile-String   : "rad-sla"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : "open-auth-vport"
Category-Map-Name    : ""
Dhcp6 ClientId (DUID): 00030001000065171001
Dhcp6 IAID           : 0
Dhcp6 IAID Type      : non-temporary
Dhcp6 Client Ip      : FE80::200:65FF:FE17:1001
Primary-Dns          : N/A
Secondary-Dns        : N/A
Pool Name            : "ludb"
Dhcp6 Server Addr    : 2001:DB8::1
Dhcp6 ServerId (DUID): 00030001d897ff000000
Dhcp6 InterfaceId    : open-dhcp
Dhcp6 RemoteId       : 0000ipoe-v6

Lease Info origin    : DHCP

ServerLeaseStart     : 07/02/2013 10:07:03
ServerLastRenew      : 07/02/2013 10:07:03
ServerLeaseEnd       : 07/02/2013 11:07:03
Session-Timeout      : N/A
Radius User-Name     : "open-dhcp@alu-domain"
-------------------------------------------------------------------------------
```

IP Proxy Case with LUDB + RADIUS Authentication

IP address is assigned via RADIUS.

- RADIUS provides IP addresses (IPv6 lease-times are provided under the group-interface) and related parameters (DNS server, IPv4 default-gateway, etc), inter-dest-id string for Vport assignment and a framed route.
- LUDB provides sub/sla-profile strings and msap-defaults (routing context parameters and msap-policy).

Vport aggregate rate limit and the port scheduler are now added to the physical port. The Vport is associated with the subscriber through the inter-dest-id string obtained via the LUDB.

```
configure  port 1/1/5
     ethernet
       mode access
     encap-type qinq
      egress-scheduler-policy "port"
      access
           egress
              vport "open-dhcp" create
              agg-rate-limit 500
                  host-match dest "open-auth-vport" create
               exit
            exit
          exit
          no shutdown
       exit
```

The LUDB is used to assign the sub/sla-profile strings.

The domain name **alu-domain** is appended to the username (circuit-id = open-dhcp or username = open-pppoe) before an Access-Request is sent to the RADIUS server that is configured in the authentication policy **open-dhcp**.

```
local-user-db "open-dhcp" create
    ipoe
        match-list circuit-id
        host "open-dhcp" create
            host-identification
                circuit-id string "open-dhcp"
            exit
            auth-policy "open-dhcp"
            auth-domain-name "alu-domain"
            identification-strings 254 create
                sla-profile-string "ludb-sla"
                sub-profile-string "ludb-sub"
            exit
            msap-defaults
                group-interface "open-auth"
                policy "msaps"
                service 10
            exit
            no shutdown
        exit
    exit
    ppp
        match-list circuit-id mac username
        host "open-ppp" create
            host-identification
                username "open-pppoe"
            exit
            auth-policy "open-dhcp"
```

```
                    password chap "ALU" hash2
                    identification-strings 254 create
                        sla-profile-string "ludb-sla"
                        sub-profile-string "ludb-sub"
                    exit
                    msap-defaults
                        group-interface "open-auth"
                        policy "msaps"
                        service 10
                    exit
                    no shutdown
                exit
            exit
            no shutdown
        exit
```

RADIUS is defined in the **authentication-policy**. The domain name can be appended to the PPPoE subscriber host directly via authentication-policy, while for IPoE subscribers the domain name is appended via authentication-policy in conjunction with LUDB.

```
configure subscriber-mgmt authentication-policy "open-dhcp"
password "ALU" hash2
      ppp-user-name append "alu-domain"
      radius-authentication-server
           server 1 address X.Y.Z.W secret "ALU" hash2
       exit
      user-name-format circuit-id append
       pppoe-access-method pap-chap
```

The RADIUS user configuration file uses the domain extension as inserted by the 7x50 BNG node to authenticate the user. The *inter-dest-id* string and the host IP address are provided by the RADIUS server (proxy case) along with other IP addressing parameters.

The IPv4 lease time (30 minutes) for IPv4 addresses are provided by the RADIUS server, while the lease time (30 minutes) for IPv6 addresses/prefixes are configured under the **group-interface**.

```
open-dhcp@alu-domain   Cleartext-Password := "ALU"
            Alc-Int-Dest-Id-Str = open-auth-vport,
            Framed-IP-Address = 10.10.0.230,
            Framed-IP-Netmask = 255.255.255.0,
            Alc-Default-Router = 10.10.0.1,
            Alc-Lease-Time = 1800,
            Client-DNS-Pri = 172.16.20.20,
            Client-DNS-Sec = 172.16.20.21,
            Alc-IPv6-Address = 2001:db8::100,
            Delegated-IPv6-Prefix = 2001:DB8:40:100::/56,
            Alc-IPv6-Primary-Dns = 2001:DB8::2000,
            Alc-Ipv6-Secondary-Dns = 2001:DB8::2001,
            Framed-Route = "192.168.1.0/24 0.0.0.0",
            Fall-Through = No
```

```
open-pppoe@alu-domain  Cleartext-Password := "ALU"
            Alc-Int-Dest-Id-Str = open-auth-vport,
            Framed-IP-Address = 10.10.0.231,
            Framed-IP-Netmask = 255.255.255.255,
            Client-DNS-Pri = 172.16.20.20,
            Client-DNS-Sec = 172.16.20.21,
            Alc-IPv6-Address = 2001:db8:0:1::100,
            Delegated-IPv6-Prefix = 2001:DB8:40:200::/56,
            Alc-IPv6-Primary-Dns = 2001:DB8::2000,
            Alc-Ipv6-Secondary-Dns = 2001:DB8::2001,
            Framed-Route = "192.168.2.0/24 0.0.0.0",
            Fall-Through = No
```

The group-interface configuration is shown below. Note that common parts of the configuration as defined earlier still apply.

```
configure service vprn 10 subscriber-interface "sub-int-1" group-interface "open-
auth" create
      ipv6
        dhcp6
           proxy-server
              renew-timer min 7
                 rebind-timer min 10
              valid-lifetime min 30
               preferred-lifetime min 15
              client-applications dhcp ppp
                 no shutdown
           exit
        exit
    exit
      dhcp
        proxy-server
              emulated-server 10.12.0.1
            no shutdown
          exit
      no shutdown
      exit
exit
```

RADIUS sub/sla-profiles supplied via the LUDB are used:

```
configure  subscriber-mgmt sla-profile "ludb-sla"
     description "sla-profile obtained via LUDB"
      host-limit 100
      egress
           qos 10 vport-scheduler
      exit
ip-filter 10
exit

config>subscr-mgmt# sub-profile "ludb-sub"
description "sub-profile obtained via LUDB"
      egress
           agg-rate-limit 15000
      exit
```

Show Commands

The following command shows that the LUDB-sub/sla-profiles, as supplied via LUDB, are in use. The IP addresses are supplied via the RADIUS server. The subscriber-id is auto-generated (not returned via LUDB or RADIUS) and it is set to circuit-id for the IPoE subscriber-host, and to the *username*|*session-id* combination for PPPoE subscriber host.

```
*A:BNG-1# show service  active-subscribers
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber open-dhcp (ludb-sub)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[1/1/5:17.10] - sla:ludb-sla
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.10.0.230
                00:00:65:17:10:01 N/A       DHCP
2001:DB8::100/128
                00:00:65:17:10:01 N/A       IPoE-DHCP6
2001:DB8:40:100::/56
                00:00:65:17:10:01 N/A       IPoE-DHCP6
-------------------------------------------------------------------------------
Subscriber open-pppoe|12 (ludb-sub)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[1/1/5:17.11] - sla:ludb-sla
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.10.0.231
                00:00:65:17:11:02 12        IPCP
2001:DB8::1:0:0:0:100/128
                00:00:65:17:11:02 12        PPP-DHCP6
2001:DB8:40:200::/56
                00:00:65:17:11:02 12        PPP-DHCP6
-------------------------------------------------------------------------------
Number of active subscribers : 2
-------------------------------------------------------------------------------
```

The following command shows more details about the subscriber-host, such as the group-interface, vport, address origin, acct-session-id, etc. Vport is selected based on the *inter-dest-id* string supplied via RADIUS.

For the purpose of brevity, the output for only two hosts is shown, one with IPv4 address and one with IPv6 prefix. The remaining IP addresses/prefixes are not shown since the output follows the same logic.

```
*A:BNG-1# show service  id 10 subscriber-hosts detail
===============================================================
Subscriber Host table
===============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID Origin       Fwding State
---------------------------------------------------------------
[1/1/5:17.10]          open-dhcp
  10.10.0.230
    00:00:65:17:10:01    N/A       DHCP         Fwding
---------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : open-auth
Sub Profile           : ludb-sub
SLA Profile           : ludb-sla
App Profile           : N/A
Egress Q-Group        : policer-output-queues
Egress Vport          : open-dhcp
Acct-Session-Id       : D897FF0000004751D31B6E
Acct-Q-Inst-Session-Id: D897FF0000004851D31B6E
Address Origin        : AAA
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A

---------------------------------------------------------------
[1/1/5:17.11]          open-pppoe|12
  2001:DB8:40:200::/56
    00:00:65:17:11:02    12        PPP-DHCP6    Fwding
---------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : open-auth
Sub Profile           : ludb-sub
SLA Profile           : ludb-sla
App Profile           : N/A
Egress Q-Group        : policer-output-queues
Egress Vport          : open-dhcp
Acct-Session-Id       : D897FF0000004651D31B6B
Acct-Q-Inst-Session-Id: D897FF0000004451D31B65
Address Origin        : AAA
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
---------------------------------------------------------------
Number of subscriber hosts : 6            The remaining 4 hosts are not shown here.
===============================================================
```

The following command shows that the subscriber identity is set to *circuit-id (plus session-id)* as instructed by the **auto-sub-id-key** command (the *subscriber-id* string is not returned via LUDB or RADIUS). The lease times are set to 30 minutes as defined by RADIUS for IPv4 addresses and by the group-interface for IPv6 addresses/prefixes (proxy-case). The username passed to RADIUS is the circuit-id or username appended to the *alu-dmain* domain name. The origin of the lease is RADIUS.

```
*A:BNG-1# show service  id 10 dhcp lease-state detail
===============================================================================
DHCP lease states for service 10
===============================================================================
Service ID           : 10
IP Address           : 10.10.0.230
Client HW Address    : 00:00:65:17:10:01
Subscriber-interface : sub-int-1
Group-interface      : open-auth
SAP                  : [1/1/5:17.10]
Up Time              : 0d 00:02:24
Remaining Lease Time : 0d 00:27:37
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "open-dhcp"
Sub-Profile-String   : "ludb-sub"
SLA-Profile-String   : "ludb-sla"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : "open-auth-vport"
Category-Map-Name    : ""

Lease Info origin    : Radius

Ip-Netmask           : 255.255.255.0
Broadcast-Ip-Addr    : 10.10.0.255
Default-Router       : 10.10.0.1
Primary-Dns          : 172.16.20.20
Secondary-Dns        : 172.16.20.21
Primary-Nbns         : N/A
Secondary-Nbns       : N/A

ServerLeaseStart     : 07/02/2013 11:26:54
ServerLastRenew      : 07/02/2013 11:26:54
ServerLeaseEnd       : 07/02/2013 11:56:54
Session-Timeout      : N/A
Lease-Time           : 0d 00:30:00
DHCP Server Addr     : N/A

Relay Agent Information
  Circuit Id         : open-dhcp
  Remote Id          : ipoe-v6
Radius User-Name     : "open-dhcp@alu-domain"

Managed Routes       : 192.168.1.0/24           installed
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
```

For the purpose of brevity, the output for only two IPv6 prefixes are shown. The remaining two IPv6 leases are not shown since the output follows the same logic.

```
*A:BNG-1# show service  id 10 dhcp6 lease-state detail
===============================================================================
DHCP lease states for service 10
===============================================================================
Service ID           : 10
```

```
IP Address           : 2001:DB8:40:100::/56
Client HW Address    : 00:00:65:17:10:01
Subscriber-interface : sub-int-1
Group-interface      : open-auth
SAP                  : [1/1/5:17.10]
Up Time              : 0d 00:02:32
Remaining Lease Time : 0d 00:27:28
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "open-dhcp"
Sub-Profile-String   : "ludb-sub"
SLA-Profile-String   : "ludb-sla"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : "open-auth-vport"
Category-Map-Name    : ""
Dhcp6 ClientId (DUID): 00030001000065171001
Dhcp6 IAID           : 0
Dhcp6 IAID Type      : prefix
Dhcp6 Client Ip      : FE80::200:65FF:FE17:1001
Primary-Dns          : 2001:DB8::2000
Secondary-Dns        : 2001:DB8::2001
Pool Name            : ""
Dhcp6 Server Addr    : N/A
Dhcp6 ServerId (DUID): N/A
Dhcp6 InterfaceId    : open-dhcp
Dhcp6 RemoteId       : 0000ipoe-v6

Lease Info origin    : Radius

ServerLeaseStart     : 07/02/2013 11:26:58
ServerLastRenew      : 07/02/2013 11:26:58
ServerLeaseEnd       : 07/02/2013 11:56:58
Session-Timeout      : N/A
Radius User-Name     : "open-dhcp@alu-domain"
-------------------------------------------------------------------------------
Service ID           : 10
IP Address           : 2001:DB8:40:200::/56
Client HW Address    : 00:00:65:17:11:02
Subscriber-interface : sub-int-1
Group-interface      : open-auth
SAP                  : [1/1/5:17.11]
Up Time              : 0d 00:02:39
Remaining Lease Time : 0d 00:27:21
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "open-pppoe|12"
Sub-Profile-String   : "ludb-sub"
SLA-Profile-String   : "ludb-sla"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : "open-auth-vport"
Category-Map-Name    : ""
Dhcp6 ClientId (DUID): 00030001000065171102
Dhcp6 IAID           : 0
Dhcp6 IAID Type      : prefix
Dhcp6 Client Ip      : FE80::200:65FF:FE17:1102
```

```
Primary-Dns          : 2001:DB8::2000
Secondary-Dns        : 2001:DB8::2001
Pool Name            : ""
Dhcp6 Server Addr    : N/A
Dhcp6 ServerId (DUID): N/A
Dhcp6 InterfaceId    : open-pppoe
Dhcp6 RemoteId       : N/A

Lease Info origin    : Radius

ServerLeaseStart     : 07/02/2013 11:26:51
ServerLastRenew      : 07/02/2013 11:26:51
ServerLeaseEnd       : 07/02/2013 11:56:51
Session-Timeout      : N/A
Radius User-Name     : "open-pppoe@alu-domain"
-------------------------------------------------------------------------------
Number of lease states : 4              The remaining two not shown in this output.
===============================================================================
```

IP Proxy Case with LUDB + RADIUS Authentication

P address is assigned via LUDB.

- RADIUS provides sub/sla-profile strings and a framed IPv4 route.
- LUDB provides IP addresses (IPv6 lease-times are provided under the group-interface) and related parameters (DNS server, IPv4 default-gateway, etc), inter-dest-id string for Vport assignment and msap-defaults (routing context parameters and msap-policy).

Vport aggregate rate limit and the port scheduler are now added to the physical port. The Vport is associated with the subscriber through the inter-dest-id string obtained via the LUDB.

```
configure  port 1/1/5
     ethernet
       mode access
     encap-type qinq
      egress-scheduler-policy "port"
      access
           egress
              vport "open-dhcp" create
              agg-rate-limit 500
                  host-match dest "open-auth-vport" create
               exit
             exit
          exit
          no shutdown
        exit
```

The LUDB is used to assign the inter-dest-id string, host IP addresses and IP addressing parameters. The DHCP lease time for IPv4 addresses is set to 15 minutes in the LUDB while lease times for IPv6 addresses/prefixes is set under the group-interface (set to 30 minutes).

Domain name *alu-domain* is appended to the username (circuit-id = *open-dhcp* or username = *open-pppoe*) before an Access-Request is sent to the RADIUS server that is configured in the authentication-policy **open-dhcp**.

```
local-user-db "open-dhcp" create
    dhcp
        match-list circuit-id
        host "open-dhcp" create
            host-identification
                circuit-id string "open-dhcp"
                exit
                address 10.10.0.230
                auth-policy "open-dhcp"
                auth-domain-name "alu-domain"
                identification-strings 254 create
                    inter-dest-id "open-auth-vport"
                exit
                msap-defaults
                    group-interface "open-auth"
                    policy "msaps"
                    service 10
                exit
                options
                    subnet-mask 255.255.255.0
                    default-router 10.10.0.254
                    dns-server 172.16.20.20 172.16.20.21
                    lease-time min 15
                exit
                options6
                    dns-server 2001:DB8::2000 2001:DB8::2001
                exit
                ipv6-address 2001:DB8::100
                ipv6-delegated-prefix 2001:DB8:40:100::/56
                no shutdown
            exit
        exit
```

RADIUS is defined in the authentication-policy. The domain name can be appended to the PPPoE subscriber host directly via authentication-policy while for IPoE subscribers, the domain name is appended via authentication policy in conjunction with LUDB.

```
configure  subscriber-mgmt authentication-policy "open-dhcp"
password "ALU" hash2
      ppp-user-name append "alu-domain"
      radius-authentication-server
           server 1 address X.Y.Z.W secret "ALU" hash2
       exit
      user-name-format circuit-id append
       pppoe-access-method pap-chap
```

The RADIUS user configuration file uses the domain extension as inserted by the 7x50 to authenticate the user.

```
open-dhcp@alu-domain   Cleartext-Password := "ALU"
            Alc-Subsc-Prof-Str = rad-sub,
            Alc-SLA-Prof-Str = rad-sla,
            Framed-Route = "192.168.1.0/24 0.0.0.0",
            Fall-Through = No

open-pppoe@alu-domain  Cleartext-Password := "ALU"
            Alc-Subsc-Prof-Str = rad-sub,
            Alc-SLA-Prof-Str = rad-sla,
            Framed-Route = "192.168.2.0/24 0.0.0.0",
            Fall-Through = No
```

The group interface configuration is shown below. Note that common parts of the configuration as defined earlier still apply.

```
configure service vprn 10 subscriber-interface "sub-int-1" group-interface "open-
auth" create
    ipv6
        dhcp6
            proxy-server
                renew-timer min 7
                rebind-timer min 10
                valid-lifetime min 30
                preferred-lifetime min 15
                client-applications dhcp ppp
                no shutdown
            exit
        exit
    exit
    dhcp
    proxy-server
        emulated-server 10.12.0.1
        no shutdown
    exit
    no shutdown
    exit
exit
```

RADIUS sub/sla-profiles supplied by RADIUS are defined as:

```
configure  subscriber-mgmt sla-profile "rad-sla"
     description "sla-profile obtained via LUDB"
      host-limit 100
      egress
           qos 10 vport-scheduler
      exit
ip-filter 10
exit

configure subscriber-mgmt sub-profile "rad-sub"
description "sub-profile obtained via LUDB"
      egress
           agg-rate-limit 15000
```

```
      exit
```

Show Commands

The following command shows that the rad-sub/sla-profiles, as provided by RADIUS, are in use. The IP addresses are provided by LUDB. The *subscriber-id* is auto-generated (not returned via the LUDB or RADIUS) and it is set to *circuit-id* for IPoE subscriber-host(s) and to *username*|*session-id* combination for PPPoE subscriber host(s).

```
*A:BNG-1# show service active-subscribers
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber open-dhcp (rad-sub)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[1/1/5:17.10] - sla:rad-sla
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.10.0.230
                00:00:65:17:10:01 N/A       DHCP
2001:DB8::100/128
                00:00:65:17:10:01 N/A       IPoE-DHCP6
2001:DB8:40:100::/56
                00:00:65:17:10:01 N/A       IPoE-DHCP6

-------------------------------------------------------------------------------
Subscriber open-pppoe|1 (rad-sub)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[1/1/5:17.11] - sla:rad-sla
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.10.0.231
                00:00:65:17:11:02 1         IPCP
2001:DB8::1:0:0:0:100/128
                00:00:65:17:11:02 1         PPP-DHCP6
2001:DB8:40:200::/56
                00:00:65:17:11:02 1         PPP-DHCP6

-------------------------------------------------------------------------------
Number of active subscribers : 2
-------------------------------------------------------------------------------
```

The following command shows more details about the subscriber-host, such as the group-interface, vport, address origin, acct-session-id, etc. Vport is selected based on the *inter-dest-id* string as supplied via RADIUS.

For the purpose of brevity, the output for only two hosts is shown, one with IPv4 address and one with IPv6 prefix. The remaining IP addresses/prefixes are not shown since the output follows the same logic.

```
*A:BNG-1# show service id 10 subscriber-hosts detail
===============================================================
Subscriber Host table
===============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID Origin       Fwding State
---------------------------------------------------------------
[1/1/5:17.10]          open-dhcp
  10.10.0.230
    00:00:65:17:10:01    N/A       DHCP         Fwding
---------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : open-auth
Sub Profile           : rad-sub
SLA Profile           : rad-sla
App Profile           : N/A
Egress Q-Group        : policer-output-queues
Egress Vport          : open-dhcp
Acct-Session-Id       : D897FF0000000051D48D5A
Acct-Q-Inst-Session-Id: D897FF0000000151D48D5A
Address Origin        : LUDB
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
---------------------------------------------------------------
 [1/1/5:17.11]          open-pppoe|1
  2001:DB8:40:200::/56
    00:00:65:17:11:02    1         PPP-DHCP6    Fwding
---------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : open-auth
Sub Profile           : rad-sub
SLA Profile           : rad-sla
App Profile           : N/A
Egress Q-Group        : policer-output-queues
Egress Vport          : open-dhcp
Acct-Session-Id       : D897FF0000000851D48D66
Acct-Q-Inst-Session-Id: D897FF0000000651D48D61
Address Origin        : LUDB
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
---------------------------------------------------------------
Number of subscriber hosts : 6            The remaining 4 hosts are not shown here.
===============================================================
```

The following command shows that the subscriber identity is set to circuit-id (plus session-id) as instructed by the **auto-sub-id-key** command (the *subscriber-id* string is not returned via the LUDB or RADIUS). The DHCPv4 lease time is set to set to 15 minutes as defined by the LUDB. The DHCPv6 lease times are set to 30 minutes as configured under the group-interface. The username passed to RADIUS is the circuit-id or username appended by the *alu-dmain* domain name. The origin of the lease is RADIUS.

```
*A:BNG-1# show service  id 10 dhcp lease-state detail
===============================================================================
DHCP lease states for service 10
===============================================================================
Service ID           : 10
IP Address           : 10.10.0.230
Client HW Address    : 00:00:65:17:10:01
Subscriber-interface : sub-int-1
Group-interface      : open-auth
SAP                  : [1/1/5:17.10]
Up Time              : 0d 00:02:09
Remaining Lease Time : 0d 00:12:51
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "open-dhcp"
Sub-Profile-String   : "rad-sub"
SLA-Profile-String   : "rad-sla"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : "open-auth-vport"
Category-Map-Name    : ""

Lease Info origin    : UserDb

Ip-Netmask           : 255.255.255.0
Broadcast-Ip-Addr    : 10.10.0.255
Default-Router       : 10.10.0.254
Primary-Dns          : 172.16.20.20
Secondary-Dns        : 172.16.20.21
Primary-Nbns         : N/A
Secondary-Nbns       : N/A

ServerLeaseStart     : 07/03/2013 13:45:14
ServerLastRenew      : 07/03/2013 13:45:14
ServerLeaseEnd       : 07/03/2013 14:00:14
Session-Timeout      : N/A
Lease-Time           : 0d 00:15:00
DHCP Server Addr     : N/A

Relay Agent Information
  Circuit Id         : open-dhcp
  Remote Id          : ipoe-v6
Radius User-Name     : "open-dhcp@alu-domain"

Managed Routes       : 192.168.1.0/24           installed
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
```

For the purpose of brevity, the output for only two IPv6 leases is shown. The remaining two IPv6 leases are not shown since the output follows the same logic.

```
*A:BNG-1# show service  id 10 dhcp6 lease-state detail
===============================================================================
DHCP lease states for service 10
===============================================================================
Service ID           : 10
IP Address           : 2001:DB8::100/128
Client HW Address    : 00:00:65:17:10:01
Subscriber-interface : sub-int-1
Group-interface      : open-auth
SAP                  : [1/1/5:17.10]
Up Time              : 0d 00:02:17
Remaining Lease Time : 0d 00:27:43
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "open-dhcp"
Sub-Profile-String   : "rad-sub"
SLA-Profile-String   : "rad-sla"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : "open-auth-vport"
Category-Map-Name    : ""
Dhcp6 ClientId (DUID): 00030001000065171001
Dhcp6 IAID           : 0
Dhcp6 IAID Type      : non-temporary
Dhcp6 Client Ip      : FE80::200:65FF:FE17:1001
Primary-Dns          : 2001:DB8::2000
Secondary-Dns        : 2001:DB8::2001
Pool Name            : ""
Dhcp6 Server Addr    : N/A
Dhcp6 ServerId (DUID): N/A
Dhcp6 InterfaceId    : open-dhcp
Dhcp6 RemoteId       : 0000ipoe-v6

Lease Info origin    : UserDb

ServerLeaseStart     : 07/03/2013 13:45:17
ServerLastRenew      : 07/03/2013 13:45:17
ServerLeaseEnd       : 07/03/2013 14:15:17
Session-Timeout      : N/A
Radius User-Name     : "open-dhcp@alu-domain"
-------------------------------------------------------------------------------
Service ID           : 10
IP Address           : 2001:DB8:40:200::/56
Client HW Address    : 00:00:65:17:11:02
Subscriber-interface : sub-int-1
Group-interface      : open-auth
SAP                  : [1/1/5:17.11]
Up Time              : 0d 00:02:09
Remaining Lease Time : 0d 00:27:51
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "open-pppoe|1"
Sub-Profile-String   : "rad-sub"
```

```
SLA-Profile-String   : "rad-sla"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : "open-auth-vport"
Category-Map-Name    : ""
Dhcp6 ClientId (DUID): 00030001000065171102
Dhcp6 IAID           : 0
Dhcp6 IAID Type      : prefix
Dhcp6 Client Ip      : FE80::200:65FF:FE17:1102
Primary-Dns          : 2001:DB8::2000
Secondary-Dns        : 2001:DB8::2001
Pool Name            : ""
Dhcp6 Server Addr    : N/A
Dhcp6 ServerId (DUID): N/A
Dhcp6 InterfaceId    : open-pppoe
Dhcp6 RemoteId       : N/A

Lease Info origin    : UserDb

ServerLeaseStart     : 07/03/2013 13:45:26
ServerLastRenew      : 07/03/2013 13:45:26
ServerLeaseEnd       : 07/03/2013 14:15:26
Session-Timeout      : N/A
Radius User-Name     : "open-pppoe@alu-domain"
-------------------------------------------------------------------------------
Number of lease states : 4            The remaining lease states are not shown here.
===============================================================================
```

Troubleshooting Commands

The following output shows the debugging commands which can be used to troubleshoot problems with the different authentication models.

```
*A:BNG-1# show debug
debug
    router "Base"
        radius
            server-address X.Y.Z.Y
            packet-type authentication
            detail-level medium
        exit
    exit
    router "10"
        ip
            dhcp
                detail-level high
                mode egr-ingr-and-dropped
            exit
            dhcp6
                mode egr-ingr-and-dropped
                detail-level high
            exit
        exit
        local-dhcp-server "local"
            detail-level high
```

```
                mode egr-ingr-and-dropped
            exit
            local-dhcp-server "v6"
                detail-level high
                mode egr-ingr-and-dropped
            exit
        exit
        mirror-source 100
            port 1/1/5 egress ingress
            no shutdown
        exit
        service
            id 2
                dhcp
                    mode egr-ingr-and-dropped
                    detail-level high
                    sap 1/1/5:17.*
                exit
                dhcp6
                    mode all
                    detail-level medium
                    sap 1/1/5:17.*
                exit
                ppp
                    packet
                        mode dropped-only
                        detail-level high
                        discovery
                        ppp
                        dhcp-client
                    exit
                exit
            exit
            id 10
                ppp
                    packet
                        mode egr-ingr-and-dropped
                        detail-level high
                        discovery
                        ppp
                        dhcp-client
                    exit
                exit
            exit
        exit
        subscriber-mgmt
            local-user-db open-dhcp
                detail all
            exit
        exit
    exit
```

Conclusion

The flexible authentication model allows access to various sources (LUDB, RADIUS, and Python) of subscriber parameters during the subscriber establishment phase. This model can be utilized for IPoE, PPPoE or L2TP subscribers in IES or VPRN services (including a wholesale/retail VRF model). A typical use case would be in a wholesale/retail environment where the wholesaler enforces its own rules via the LUDB before it passes the authentication request to the retailer’s RADIUS server.

Ingress Multicast Path Management

This chapter provides information about Ingress Multicast Path Management (IMPM).

Topics in this chapter include:

Applicability

The information in this chapter is applicable to all 7x50 platforms configured with IOM1/2, IOM3-XP and IMMs line cards, but not to the SR-1, ESS-1, 7710 SR c-4/12 or the 7750 SR c-4/12. The configuration was tested on release 9.0R6. There are no pre-requisites for this configuration.

Summary

Ingress Multicast Path Management (IMPM) optimizes the IPv4 and IPv6 multicast capacity on the applicable systems with the goal of achieving the maximum system-wide IP multicast throughput. It controls the delivery of IPv4/IPv6 routed multicast groups and of VPLS (IGMP and PIM) snooped IPv4 multicast groups, which usually relate to the distribution of IP TV channels.

A description is also included of the use of IMPM resources by point-to-multipoint LSP IP multicast traffic, and policed ingress routed IP multicast or VPLS broadcast, unknown or multicast traffic. The system capacity for these traffic types can be increased even with IMPM disabled.

Overview

IMPM introduces the concept of paths on a line card (IOM/IMM) which connect to planes on the chassis switch fabric (Figure 150).

IMPM monitors the ingress rate of IP multicast channels (S,G multicast streams) on line card paths and optimizes the use of the capacity of each switch fabric plane. Its goal is to forward as many channels as possible through the system in order to make maximum use of the switch fabric planes without incurring multicast packet loss. IMPM achieves this by moving entire multicast channels between the line card paths, and therefore between switch fabric planes, to achieve an optimal packing of channels onto path/planes. These actions take into consideration the total ingress multicast traffic being received by all line cards with IMPM enabled and a configured preference of each channel.

Note: S,G refers to an individual multicast stream by referencing the source (S) and multicast group (G) used by the steam.

Figure 150 **IOM/IMM Paths Connecting to Switch Fabric Planes**

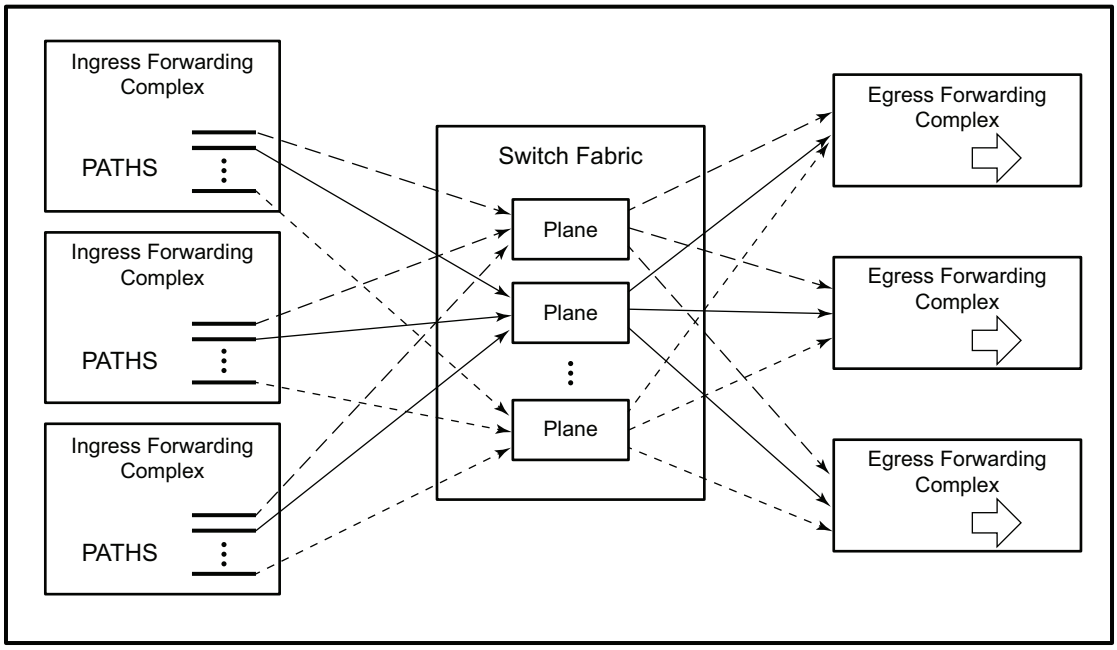

There are three types of path: primary, secondary and ancillary paths (the ancillary path is specific to the IOM1/2 and is discussed in Ancillary Path).

When a new channel is received on a line card for which there is an egress join, its traffic is initially placed on to a secondary path by default. IMPM monitors the channel's traffic rate and, after an initial monitoring period, can move the channel to another path (usually a primary path) on which sufficient capacity exists for the channel. IMPM constantly monitors all of the ingress channels and therefore keeps a picture of the current usage of all the line card paths and switch fabric planes. As new channels arrive, IMPM assigns them onto available capacity, which may involve moving existing channels between paths (and planes). If a channel stops, IMPM will be aware that more capacity is now available. If the traffic rate of any channel(s) changes, IMPM will continue to optimize the use of the path/planes.

In the case where there is insufficient plane capacity available for all ingress channels, entire channel(s) are blackholed (dropped) rather than allowing a random degradation across all channels. This action is based on a configurable channel preference with the lowest preference channel being dropped first. If path/plane capacity becomes available, then the blackholed channel(s) can be re-instated.

Paths and Planes

Each path connects to one plane on the switch fabric which is then used to replicate multicast traffic to the egress line cards. Further replication can occur on the egress line card but this is not related to IMPM so is not discussed.

Each plane has a physical connection to every line card and operates in full duplex mode allowing the possibility for traffic received on a plane from one line card path to be sent to every other line card, and back to the ingress line card (note that traffic is actually only sent to the line cards where it may exit). Therefore a given plane interconnects all line cards which allow ingress multipoint traffic from a line card with a path connected to this plane to be sent to multiple egress line cards.

Traffic could be sent by only one line card, or by multiple line cards simultaneously, on to a given plane. The total amount of traffic on a path or plane cannot exceed the capacity of that path or plane, respectively.

There could be more planes available on the switch fabrics than paths on the line cards. Conversely, there could be more total line card paths than planes available on the switch fabrics. In the latter case, the system distributes the paths as equally as possible over the available planes and multiple paths would be assigned to a given0 plane. Note that multiple paths of either type (primary or secondary) can terminate on a given plane.

The number of paths available per line card depends on the type of line card used whereas the number of planes on a system depends on the chassis type, the chassis mode (**a**, **b**, **c**, **d**) and the number of installed switch fabrics.

To clarify these concepts, consider a system with the following hardware installed.

```
A:PE-1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned      Equipped         Admin   Operational        Comments
          Card-type        Card-type        State   State
-------------------------------------------------------------------------------
6         iom3-xp          iom3-xp          up      up
7         imm8-10gb-xfp    imm8-10gb-xfp    up      up
8         iom3-xp          iom3-xp          up      up
A         sfm4-12          sfm4-12          up      up/active
B         sfm4-12          sfm4-12          up      up/standby
===============================================================================
A:PE-1#
```

Output 1 shows the mapping of paths to switch fabric planes.

```
A:PE-1# show system switch-fabric high-bandwidth-multicast
===============================================================================
Switch Fabric
===============================================================================
         Cap:               Planes:
Slot/Mda Min  Max  Hbm Grp  Hi | Lo
-------------------------------------------------------------------------------
 6/1     100% 100% No  0     1  0  3  4  5  6  7  8  9 10 11 12 13 14 15 | 16
 7/1     100% 100% No  0    19 17 20 21 22 23 24 25 26 27 28 29 30 31 32 | 33
 8/1     100% 100% No  0    35 34 36 37 38 39 40 41 42 43 44 45 46 47  0 |  1
 A       100% 100% No  0     2 |  2
 B       100% 100% No  0     2 |  2
===============================================================================
A:PE-1#
```

Output 1: Paths and Planes in Chassis Mode d

This system has two SF/CPM4s and is using chassis mode **d**, this creates 24 planes per SF/CPM4 to give a total of 48 planes which are numbered 0-47. The IOM3-XP/IMMs have 16 paths each which are connected to different planes. The SF/CPM4s together use a single plane and an additional plane (18, which is not in the output above) is used by the system itself. As there are more paths (3x16=48) in this configuration than available planes (48-2[system planes 2,18]=46), some planes are shared by multiple paths, namely planes 0 and 1. Note that the path to plane mapping can change after a reboot or after changing hardware.

The following output shows the equivalent information if an IOM2 is added to this configuration in slot 5. In order for the IOM2 to be recognized, the system must be changed to use chassis mode **a**, **b** or **c**.

```
A:PE-1# show card
===============================================================================
Card Summary
```

```
===============================================================================
Slot      Provisioned      Equipped         Admin   Operational        Comments
          Card-type        Card-type        State   State
-------------------------------------------------------------------------------
5         iom2-20g         iom2-20g         up      up
6         iom3-xp          iom3-xp          up      up
7         imm8-10gb-xfp    imm8-10gb-xfp    up      up
8         iom3-xp          iom3-xp          up      up
A         sfm4-12          sfm4-12          up      up/active
B         sfm4-12          sfm4-12          up      up/standby
===============================================================================
A:PE-1#
```

The following output shows the mapping of the line card paths to the switch fabric planes with the IOM2 installed.

```
A:PE-1# show system switch-fabric high-bandwidth-multicast
===============================================================================
Switch Fabric
===============================================================================
         Cap:               Planes:
Slot/Mda Min  Max  Hbm Grp  Hi | Lo
-------------------------------------------------------------------------------
 5/1     100% 100% No  0     1 |  0
 5/2     100% 100% No  0     4 |  3
 6/1     100% 100% No  0     6  5  7  8  9 10 11 12 13 14 15  0  1  3  4 |  5
 7/1     100% 100% No  0     7  6  8  9 10 11 12 13 14 15  0  1  3  4  5 |  6
 8/1     100% 100% No  0     8  7  9 10 11 12 13 14 15  0  1  3  4  5  6 |  7
 A       100% 100% No  0     2 |  2
 B       100% 100% No  0     2 |  2
===============================================================================
A:PE-1#
```

Output 2: Paths and Planes in Chassis Mode a/b/c

Now that the system is not in chassis mode **d**, in fact it is in mode a (but the output would be the same in modes **b** or **c**) the SF/CPM4s each create 8 planes giving a total of 16, numbered 0-15. One plane (2) is used by the SF/CPM4s, leaving 15 (0-1,3-15) planes for connectivity to the line card paths. Each IOM2 forwarding complex has 2 paths, so the paths of the IOM2 in slot 5 are using planes 0 and 1, and 3 and 4. Note that there are now fewer planes available and more paths, so there is more sharing of planes between paths than when chassis mode **d** was used.

IMPM Managed Traffic

IMPM manages IPv4/IPv6 routed multicast traffic and VPLS (IGMP and PIM) snooped IPv4 multicast traffic, traffic that matches a <*,G> or a <S,G> multicast record in the ingress forwarding table. It manages IP multicast traffic on a bud LSR when using point-to-multipoint (P2MP) LSPs but it does not manage IP protocol control traffic or traffic using multipoint-shared queuing. Traffic being managed by IMPM involves IMPM monitoring and potentially moving the related channels between paths/planes. The unmanaged traffic rates are also monitored and taken into account in the IMPM algorithm.

Care should be taken when using the mrouter-port configuration in a VPLS service. This creates a (*,*) multicast record and consequently all multicast channels that are not delivered locally to a non-mrouter port will be treated by IMPM as a single channel.

Configuration

This section covers:

- IMPM on an IOM3-XP/IMM
- IMPM on an IOM1/2
- IMPM Not Enabled

Prerequisites

As IMPM operates on IPv4/IPv6 routed or VPLS IGMP/PIM snooped IPv4 multicast traffic, some basic multicast configuration must be enabled. This section uses routed IP multicast in the global routing table which requires IP interfaces to be configured with PIM and IGMP. The configuration uses a PIM rendezvous point and static IGMP joins. The following is an example of the complete configuration of one interface.

```
configure
    router
        interface "int-IOM3-1"
            address 172.16.6.254/24
            port 6/2/1
        exit
        igmp
            interface "int-IOM3-1"
                static
                    group 239.255.0.1
```

```
                    starg
                exit
            exit
        exit
        no shutdown
    exit
    pim
        interface "int-IOM3-1"
        exit
        rp
            static
                address 192.0.2.1
                    group-prefix 239.255.0.0/16
                exit
            exit
        exit
        no shutdown
    exit
exit
```

One interface is configured on each line card configured in the system, as shown in the following output, but omitting their IGMP and PIM configuration.

```
configure
    router
        interface "int-IMM8"
            address 172.16.3.254/24
            port 7/2/1
        exit
        interface "int-IOM2"
            address 172.16.1.254/24
            port 5/2/1
        exit
        interface "int-IOM3-1"
            address 172.16.2.254/24
            port 6/2/1
        exit
        interface "int-IOM3-2"
            address 172.16.4.254/24
            port 8/2/1
        exit
    exit
```

Configuring IMPM

The majority of the IMPM configuration is performed under the **mcast-management** CLI nodes and consists of:

a. The bandwidth-policy for characteristics relating to the IOM/IMM paths. This is applied on an IOM3-XP/IMM fp (fp is the system term for a forwarding complex on an IOM3-XP/IMM.), or an IOM1/2 MDA, under ingress mcast-path-management, with a bandwidth-policy named default being applied by default.

– IOM1/2

```
config# card slot-number mda mda-slot
    ingress
        mcast-path-management
            bandwidth-policy policy-name
```

– IOM3-XP/IMM

```
config# card slot-number fp [1]
    ingress
        mcast-path-management
            bandwidth-policy policy-name
```

b. The multicast-info policy for information related to the channels and how they are handled by the system. To facilitate provisioning, parameters can be configured under a three level hierarchy with each level overriding the configuration of its predecessor:

– Bundle: a group of channels

– Channel: a single channel or a non-overlapping range of channels

– Source-override: channels from a specific sender

```
config# mcast-management multicast-info-policy policy-name [create]
    bundle bundle-name [create]
        channel ip-address [ip-address] [create]
            source-override ip-address [create]
```

This policy is applied where the channel enters the system, so under router or service (vpls or vprn); the latter allows the handling of channels to be specific to a service, even if multiple services use overlapping channel addresses.

```
config# router multicast-info-policy policy-name

config# service vpls service-id multicast-info-policy policy-name

config# service vprn service-id multicast-info-policy policy-name
```

A default multicast-info-policy is applied to the above when IMPM is enabled.

c. The chassis-level node configures the information relating to the switch fabric planes.

```
config# mcast-management chassis-level
```

In addition, the command hi-bw-mcast-src (under an IOM3-XP/IMM fp or an IOM1/2 MDA) can be used to control the path to plane mapping among forwarding complexes.

IMPM on an IOM3-XP/IMM

IMPM is enabled on IOM3-XP/IMMs on under the card/fp CLI node as follows

```
config# card slot-number fp 1 ingress mcast-path-management no shutdown
```

IOM3-XP/IMM Paths

16 paths are available on an IOM3-XP/IMM when IMPM is enabled which can be either primary paths or secondary paths. By default the 16 paths are divided into 15 primary paths and 1 secondary path, as can be seen using the following command with IMPM enabled only on slot 6 (this corresponds to the plane assignment in Output 1):

```
*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[6][0]: 0xf33b0a00
PATH:                                          PLANE:
Type SGs     InUseBW   AvailBW   TotalBw  ID   SGs   InUseBW   AvailBW   TotalBw
P      1           0         -         -   1     1         0   2000000   2000000
P      1           0         -         -   0     1         0   2000000   2000000
P      1           0         -         -   3     1         0   2000000   2000000
P      1           0         -         -   4     1         0   2000000   2000000
P      1           0         -         -   5     1         0   2000000   2000000
P      1           0         -         -   6     1         0   2000000   2000000
P      1           0         -         -   7     1         0   2000000   2000000
P      1           0         -         -   8     1         0   2000000   2000000
P      1           0         -         -   9     1         0   2000000   2000000
P      1           0         -         -  10     1         0   2000000   2000000
P      1           0         -         -  11     1         0   2000000   2000000
P      1           0         -         -  12     1         0   2000000   2000000
P      1           0         -         -  13     1         0   2000000   2000000
P      1           0         -         -  14     1         0   2000000   2000000
P      1           0         -         -  15     1         0   2000000   2000000
S      1           0         -         -  16     1         0   1800000   1800000
B      0           0         -         -   -     -         -         -         -
*A:PE-1#
```

Output 3: Paths/Planes on IOM3-XP/IMM

The left side of the output displays information about the paths (type {P=primary, s=secondary or B=blackholed}, number of “S,G”s and bandwidth in use (the path bandwidth cannot be set on an IOM3-XP/IMM, so the path available and total bandwidth always shows “-“)) and the right side displays similar information about the associated planes (this will be a combination of the information for all paths connected to this plane). Note that one SG is always present on each path; this is used by the system and relates to the unmanaged traffic.

The primary/secondary paths are also highlighted in the planes section of Output 1, the primary paths being connected to the planes on the left of the “|” and the secondary paths to its right. There is a default primary path and a default secondary path; these correspond to the left-most plane and right-most plane for each line card, respectively.

Primary paths are used by:

- Expedited IES, VPLS and VPRN service ingress non-managed multipoint traffic (using the SAP based queues). This uses the default primary path.
- Expedited network ingress non-managed multipoint traffic (using the network interface queues). This uses the default primary path.
- 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)
- Non-managed P2MP LSP IP multicast traffic. This does not require IMPM to be enabled, so is discussed later in IMPM Not Enabled.
- Non-managed expedited ingress policed multipoint traffic. This does not require IMPM to be enabled, so is discussed in IMPM Not Enabled.

Secondary paths are used by:

- Best-Effort IES, VPLS and VPRN service ingress non-managed multipoint traffic (using the SAP based queues). This uses the default secondary path.
- Best-Effort network ingress non-managed traffic (using the network interface multipoint queues). This uses the default secondary path.
- 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)
- Non-managed best-effort ingress policed multipoint traffic. This does not require IMPM to be enabled, so is discussed in IMPM Not Enabled.

When IMPM is enabled, the managed traffic does not use the standard multipoint queues but instead is placed onto a separate set of shared queues which are associated with the 16 paths. These queues are instantiated in an access ingress pool (called MC Path Mgmt, see “show output” section) which exists by default (this pool can be used even when IMPM is not enabled – see section “IMPM not enabled”). Statistics relating to traffic on these queues are reflected back to the standard ingress multipoint queues for accounting and troubleshooting purposes. Note that non-managed traffic continues to use the standard ingress multipoint queues, with the exception of P2MP LSP IP multicast traffic and policed multipoint traffic.

The size of the pool by default is 10% of the total ingress pool size, the reserved CBS is 50% of the pool and the default slope policy is applied. Care should be taken when changing the size of this pool as this would affect the size of other ingress pools on the line card.

```
config# mcast-management bandwidth-policy policy-name [create]
    mcast-pool percent-of-total percent-of-buffers
               resv-cbs percent-of-pool
               slope-policy policy-name
```

It is possible to configure the parameters for the queues associated with both the primary and secondary paths, and also the number of secondary paths available, within the bandwidth-policy.

```
config# mcast-management bandwidth-policy policy-name create
    t2-paths
        primary-path
            queue-parameters
                cbs percentage
                hi-priority-only percent-of-mbs
                mbs percentage
        secondary-path
            number-paths number-of-paths [dual-sfm number-of-paths]
            queue-parameters
                cbs percentage
                hi-priority-only percent-of-mbs
                mbs percentage
```

The number of primary paths is 16 minus the number of secondary paths (at least one of each type must exist). The number-paths parameter 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.

Packets are scheduled out of the path/multicast queue as follows:

- Traffic sent on primary paths is scheduled at multicast high priority while that on secondary paths is scheduled at multicast low priority.

- For managed traffic, the standard ingress forwarding class/prioritization is not used, instead IMPM managed traffic prioritization is based on a channel's preference (described in Channel Prioritization and Blackholing Control). Egress scheduling is unchanged.

Congestion handling (packet acceptance into the path/multicast queue):

- For non-managed traffic, this is based on the standard mechanism, namely the packet's enqueuing priority is used to determine whether the packet is accepted into the path multipoint queue depending on the queue mbs/cbs and the pool shared-buffers/reserved-buffers/WRED.

For managed traffic, the congestion handling is based upon the channel's preference (described later) and the channel's cong-priority-threshold which is configured in the multicast-info-policy (here under a bundle).

```
config# mcast-management multicast-info-policy policy-name [create]
    bundle bundle-name [create]
        cong-priority-threshold preference-level
```

When the preference of a channel is lower than the cong-priority-threshold setting, the traffic is treated as low enqueuing priority, when it is equal to or higher than the cong-priority-threshold it is treated as high enqueuing priority. The default cong-priority-threshold is 4.

IOM3-XP/IMM Planes

The capacity per plane for managed traffic is by default 2Gbps for a primary path and 1.8Gbps for a secondary path. The logic behind a reduced default on the secondary is to leave capacity for new streams in case the default secondary is fully used by managed streams.

The plane capacities can be configured as follows, note that this command configures the plane bandwidth associated with primary/secondary paths as seen by each line card, the TotalBw on the right side of Output 3:

```
config# mcast-management chassis-level
   per-mcast-plane-limit megabits-per-second [secondary megabits-per-second]
                [dual-sfm megabits-per-second [secondary-dual-sfm megabits-per-
second]]
```

The first parameter defines the capacity for a primary path, the second a secondary path and the dual-sfm configures these capacities when two switch fabrics are active. The maximum plane capacity is 4Gbps but for the release used here it should only be configured on 7750 SR-12 or 7450 ESS-12 systems populated with SF/CPM4(s) and 100G FP IMMs; for all other hardware combinations the maximum should be 2Gbps. Note that secondary plane capacity cannot be higher than that of the primary plane.

These values can be tuned to constrain the amount of managed multicast traffic in favour of non-managed multicast and unicast traffic.

On the IOM3-XP/IMM line cards there is no separate control of the line card path capacity, the capacity is only constrained by the plane.

IOM3-XP/IMM Path to Plane Mapping

By default all fps (line cards for IOM3-XP/IMM) are configured into the default (zero) group as seen in Output 1 and the system distributes the paths as equally as possible over the available planes. This default works well if there is a low volume of multicast traffic (compared to the plane capacity), or if there is a higher volume multicast entering only one line card where the ingress capacity does not exceed that provided by the planes the line card is connected to.

If there are more paths than planes and, for example, there is a high bandwidth multicast channel entering two different line cards it could happen that both line cards select the same plane for two paths that are used. This would result in one of the channels being blackholed if the plane capacity is exceeded, effectively reducing the available multicast capacity from that line card. In order to avoid this situation, it is possible to configure the paths in to different groups and the system will attempt to use different planes for each group.

Output 1 and Output 2 show examples of how the paths are mapped to planes.

In both cases there are more paths than planes so some planes are shared by multiple paths. The following command provides control of this mapping.

```
config# card slot-number fp [1]
    hi-bw-mcast-src [alarm] [group group-id] [default-paths-only]
```

If an fp is configured into a non-zero group (range: 1 to 32), the system will attempt to assign dedicate planes to its paths compared to other line cards in different non-zero groups. This action is dependent on there being sufficient planes available. If two line cards are assigned to the same group, they will be assigned the same planes. The **default-paths-only** parameter performs the assignment optimization only for the default primary and secondary paths and is only applicable to IOM3-XP/IMMs. The alarm keyword causes an alarm to be generated if some planes are still shared with fps in a different group.

An example of the use of this command is shown later.

Note: When VPLS IGMP and PIM snooped traffic is forwarded to a spoke or mesh SDP, by default it is sent by the switch fabric to all line card forwarding complexes on which there is a network IP interface. This is due to the dynamic nature of the way that a spoke or mesh SDP is associated with one or more egress network IP interfaces. If there is an active spoke/mesh SDP for the VPLS service on the egress forwarding complex, the traffic will be flooded on that spoke/mesh SDP, otherwise it will be dropped on the egress forwarding complex. This can be optimized by configuring an inclusion list under the spoke or mesh SDP defining which MDAs this traffic should be flooded to.

```
config>service>vpls# [spoke-sdp|mesh-sdp] sdp-id:vc-id egress
    mfib-allowed-mda-destinations
        [no] mda mda-id
```

The switch fabric flooding domain for this spoke or mesh SDP is made up only of the MDAs that have been added to the list. An empty list implies the default behavior.

It is important to ensure that the spoke or mesh SDP can only be established across the MDAs listed, for example by using RSVP with an explicit path.

IMPM Operation on IOM3-XP/IMM

This section covers:

- Principle of Operation
- Monitoring Traffic Rates
- Channel Prioritization and Blackholing Control

Principle of Operation

Where IMPM is enabled, it constantly monitors the line cards for ingress managed traffic.

When a new channel arrives it will be placed by default on to the default secondary path. IMPM determines the ingress point for the channel and then monitors the traffic of the channel within its monitoring period in order to calculate the rate of the channel. The system then searches the multicast paths/planes attached to the line card for available bandwidth. If there is sufficient capacity on such a path/plane, the channel is moved to that plane. Planes corresponding to primary paths are used first, when there is no capacity available on any primary path/plane a secondary path/ plane is used (unless the channel is explicitly configured onto a specific path type – see the following description).

If the required bandwidth is unavailable, the system will then look for any channels that ingress this or other line cards that could be moved to a different multicast plane in order to free up capacity for the new channel. Any channel that is currently mapped to a multicast plane available to the ingress line card is eligible to be moved to a different multicast plane.

If an eligible existing channel is found, whether on this or another line card, that existing channel is moved without packet loss to a new multicast plane. If necessary, this process can be repeated resulting in multiple channels being moved. The new multicast channel is then mapped to the multicast plane previously occupied by the moved channels, again this normally is using a primary path.

If no movable channel is found, then lower preference channel(s) on any ingress line card that share multicast planes with the ingress line card of the new channel can be blackholed to free up capacity for the new channel. It is also possible to both blackhole some channels and move other channels in order to free up the required capacity. If no lower preference channel is found and no suitable channel moves are possible, the new channel will be blackholed.

If required, channels can be explicitly configured to be on either a primary or secondary path. This can be done for a bundle of channels, for example

```
config# mcast-management multicast-info-policy policy-name [create]
    bundle bundle-name [create]
        explicit-sf-path {primary|secondary|ancillary}
```

Note that the ancillary path is not applicable to the IOM3-XP/IMM line cards, however, it is discussed in the section relating to the IOM1/2. If a channel on an IOM3-XP/IMM is configured onto the ancillary path it will use a primary path instead.

One secondary path on an IOM3-XP/IMM is used as a default startup path for new incoming channels. If a large amount of new channel traffic could be received within the monitoring period, it is possible that the plane associated with the default secondary path is over loaded before IMPM has time to monitor the channels' traffic rate and move the channels to a primary path (and a plane with available capacity). This can be alleviated by configuring the following command:

```
config# mcast-management chassis-level
    round-robin-inactive-records
```

When round-robin-inactive-records is enabled, the system redistributes new channels (which are referenced by inactive S,G records) among all available line card multicast (primary, secondary) paths and their switch fabric planes.

Monitoring Traffic Rates

The monitored traffic rate is the averaged traffic rate measured over a monitoring period. The monitoring period used depends on the total number of channels seen by IMPM, the minimum is a 1 second interval and the maximum a 10 seconds interval.

The way in which the system reacts to the measured rate can be tuned using the following command:

```
config# mcast-management multicast-info-policy policy-name [create]
    bundle bundle-name [create]
        bw-activity {use-admin-bw|dynamic [falling-delay seconds]}
                    [black-hole-rate kbps]
```

The default is to use the dynamic bandwidth rate, in which case a channel's active traffic rate is determined based on the measured monitored rates. IMPM then makes a decision of how to process the channel as follows.

If the channel was un-managed, IMPM will attempt to place the channel on a path/plane with sufficient available bandwidth.

If the channel was already managed, IMPM determines the highest monitored traffic rate (within a given monitoring period) in a sliding window defined by the falling-delay. This highest chosen monitored rate is then used to re-assess the placement of the channel on the path/planes; this may cause IMPM to move the channel. This mechanism prevents the active rate for a channel being reduced due to a momentarily drop in traffic rate. The default value for falling-delay is 30 seconds, with a range of 10-3600 seconds.

The above logic is shown in Figure 151 (for simplicity, the falling-delay is exactly twice the monitoring period). It can be seen that the active rate used when the traffic rate decreases follows the highest monitored rate in any falling-delay period.

Figure 151 **Dynamic Bandwidth Rate Management**

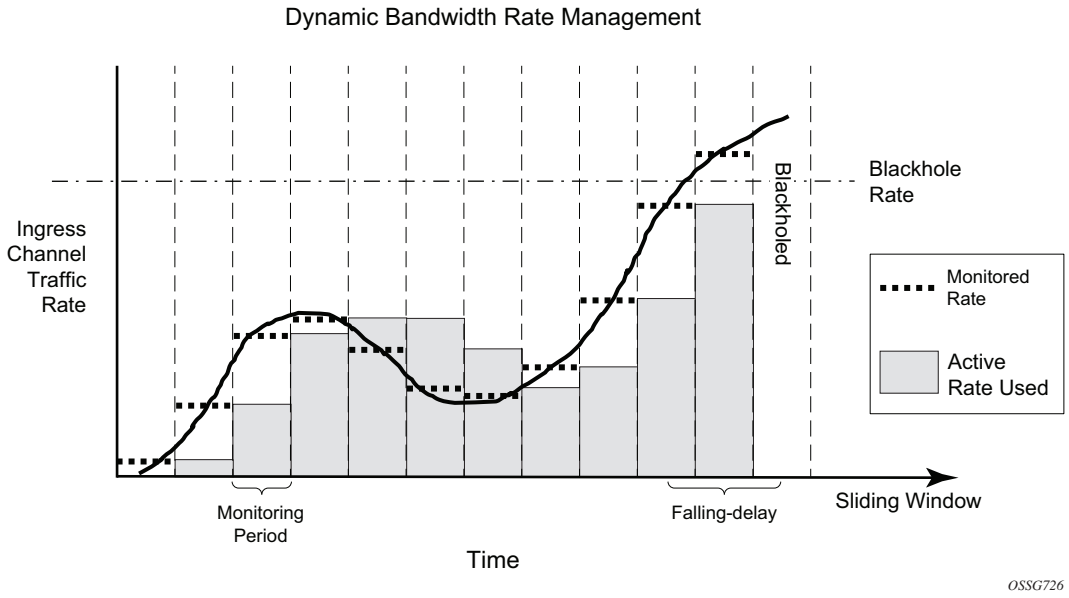

By using the sliding window of monitored rate measurements in the dynamic bandwidth measurement mode, IMPM delays releasing capacity for a channel in its calculations when the channel's rate has been reduced. This allows IMPM to ignore temporary fluctuations in a channel's rate. It is possible to tune this for cases where the reduction in a channel's rate is large by using the falling-percent-reset parameter. The default for the falling-percent-reset is 50%. Setting this to 100% effectively disables it.

```
config# mcast-management bandwidth-policy policy-name create
    falling-percent-reset percent-of-highest
```

When the monitored rate falls by a percentage which is greater or equal to falling-percent-reset, the rate used by IMPM is immediately set to this new monitored rate. This allows IMPM to react faster to significant reductions in a channel's rate while at the same time avoiding too frequent reallocations due to normal rate fluctuations. An example of the falling-percent-reset is shown in Figure 152. In the last two monitoring periods, it can be seen that the active rate used is equal to the monitored rate in the previous periods, and not the higher rate in the previous falling-delay window.

Figure 152 **Falling-Percent-Reset**

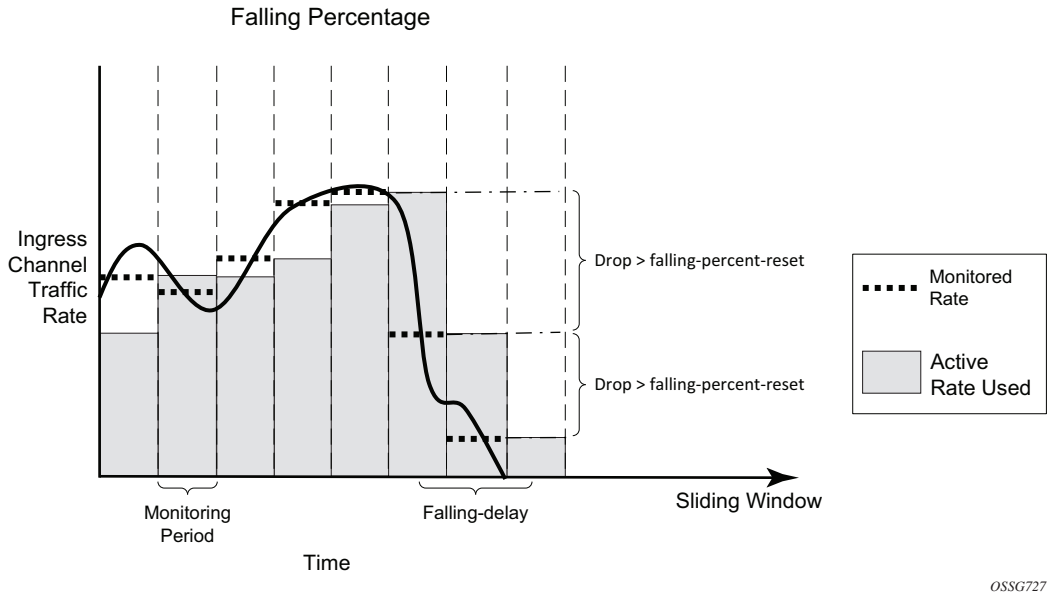

The rate management can be further tuned based on the expectation that the channel bandwidth will fluctuate around a given rate. When the bw-activity is set to use-admin-bw within the multicast-info-policy, the following parameters come into play.

```
config# mcast-management multicast-info-policy policy-name [create]
    bundle bundle-name [create]
        admin-bw kbps

config# mcast-management bandwidth-policy policy-name create
    admin-bw-threshold kilo-bits-per-second
```

IMPM will use the rate configured for the admin-bw if the monitored rate is above the admin-bw-threshold but below or equal to the admin-bw in the sliding window of the falling-delay. Whenever the monitored rate is below the admin-bw-threshold or above the admin-bw, IMPM uses the dynamic rate management mechanism. The admin-bw-threshold needs to be smaller than the admin-bw, with the latter being non-zero. This is shown in Figure 153 (for simplicity, the falling-delay is exactly twice the monitoring period). It can be seen that while the monitored rate stays between the admin-bw-threshold and the admin-bw, the active rate used is set to the admin-bw.

Figure 153 **Admin-Bw Rate Management**

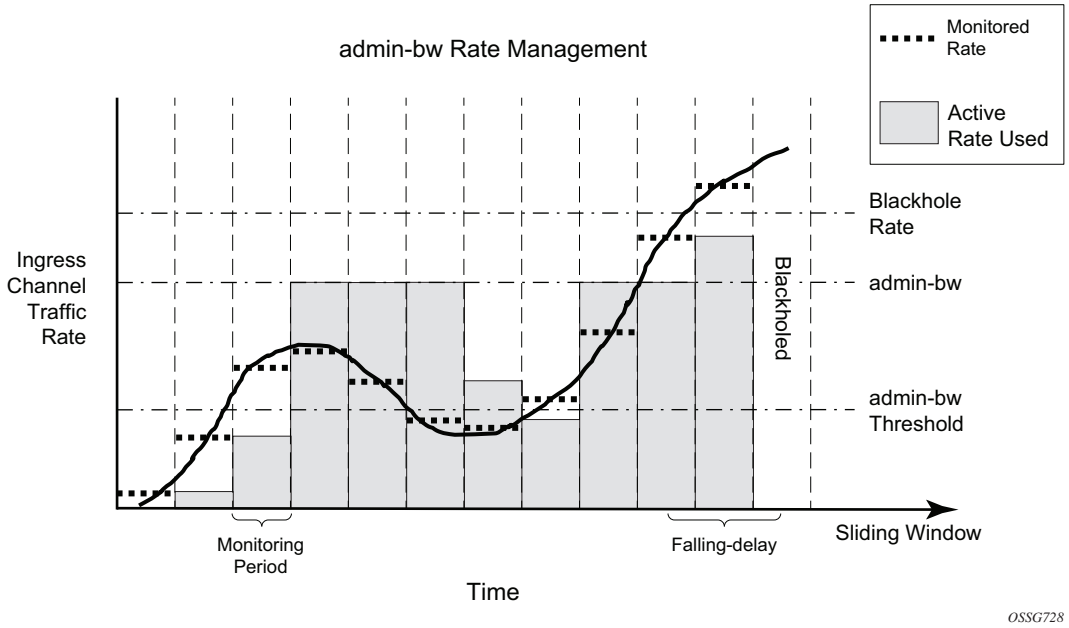

Finally, IMPM also takes into consideration the unmanaged traffic rate on the primary and secondary paths associated with SAP/network interface queues when determining the available capacity on these paths/planes. This is achieved by constantly monitoring the rate of this traffic on these queues and including this rate in the path/plane capacity usage. IMPM must be enabled on the ingress card of the unmanaged traffic otherwise it will not be monitored.

Channel Prioritization and Blackholing Control

IMPM decides which channels will be forwarded and which will be dropped based on a configured channel preference. The preference value can be in the range 0-7, with 7 being the most preferred and the default value being 0.

When there is insufficient capacity on the paths/planes to support all ingress multipoint traffic, IMPM uses the channel preferences to determine which channels should be forwarded and which should be blackholed (dropped).

This is an important distinction compared to the standard forwarding class packet prioritization; by using a channel preference, an entire channel is either forwarded or blackholed, this allows IMPM to avoid congestion having a small impact on multiple channels at the cost of entire channels being lost.

The channel preference is set within the multicast-info-policy, for example at the bundle level, with the settable values being 1-7:

```
config# mcast-management multicast-info-policy policy-name [create]
    bundle bundle-name [create]
        preference preference-level
```

The channel preference is also used for congestion control in the line card path queues – see “congestion handling” in the section on “IOM3-XP/IMM Paths” above.

Blackhole protection can also be enabled using the bw-activity command, shown above in the “Monitoring traffic rates” section. Regardless of which rate monitoring mechanism is used, a channel can be blackholed if the monitored rate exceeds the black-hole-rate, in which case the channel will be put immediately on the blackhole list and its packets dropped at ingress. This channel will no longer consume line card path or switch fabric capacity. The intention of this parameter is to provide a protection mechanism in case channels consume more bandwidth than expected which could adversely affect other channels.

The black-hole-rate can range from 0 to 40000000kbps, with no black-hole-rate by default. This protection is shown in the last monitoring period of both Figure 2 and Figure 4. Note that it will take a falling-delay period in which the channel’s rate is always below the black-hole-rate in order for the channel to be re-instated unless the reduction in the rate is above the falling-percent-reset.

IMPM on an IOM1/2

As most of the principles when using IMPM on an IOM1/2 compared to on an IOM3-XP/IMM are the same and are described above, this section focuses only on the difference between the two.

Note that an IOM1 and IOM2 have two independent 10G forwarding complexes; in both cases there is a single MDA per forwarding complex, consequently some aspects of IMPM are configured under the mda CLI node.

IMPM is enabled on an IOM1/2 under the MDA CLI node as follows:

```
config# card slot-number mda mda-slot ingress mcast-path-management no shutdown
```

IOM1/2 Paths

Each forwarding complex has three paths: one primary and one secondary path, and another type called the ancillary path which is IOM1/2 specific. The paths can be seen using the following command with IMPM enabled only on slot 5 MDA 1 and MDA 2 (referenced as [0] and [1] respectively):

```
A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[5][0]: 0xf33b0a00
PATH:                                       PLANE:
Type SGs     InUseBW   AvailBW   TotalBw  ID   SGs   InUseBW   AvailBW   TotalBw
P      1           0   2000000   2000000   1     1         0   2000000   2000000
S      1           0   1500000   1500000   0     1         0   1800000   1800000
A      0           0   5000000   5000000   -     -         -         -         -
B      0           0         -         -   -     -         -         -         -
McPathMgr[5][1]: 0xf33b3198
PATH:                                       PLANE:
Type SGs     InUseBW   AvailBW   TotalBw  ID   SGs   InUseBW   AvailBW   TotalBw
P      1           0   2000000   2000000   4     1         0   2000000   2000000
S      1           0   1500000   1500000   3     1         0   1800000   1800000
A      0           0   5000000   5000000   -     -         -         -         -
B      0           0         -         -   -     -         -         -         -
A:PE-1#
```

Output 4: Paths/Planes on IOM1/2

The primary and secondary paths function as on the IOM3-XP/IMM, specifically for:

- Traffic usage.
- Associated queues instantiated in the ingress “MC Path Mgmt” ingress pool.
- Packet scheduling.
- Congestion handling.

The queue parameters can be configured within the bandwidth-policy in a similar way to the IOM3-XP/IMM (note that the IOM3-XP/IMM equivalent for this is under the t2-paths CLI node). The bandwidth-policy is then applied under the respective MDA.

```
config# mcast-management bandwidth-policy policy-name create
    primary-path
        queue-parameters
            cbs percentage
            hi-priority-only percent-of-mbs
            mbs percentage
    secondary-path
        queue-parameters
            cbs percentage
            hi-priority-only percent-of-mbs
            mbs percentage
```

The IOM1/2 allows capacity control on the paths themselves, which is not possible on the IOM3-XP/IMM. This is achieved using the following commands.

```
config# mcast-management bandwidth-policy policy-name create
    primary-path
        path-limit megabits-per-second
    secondary-path
        path-limit megabits-per-second
```

The maximum path-limit for both the primary and secondary path is 2Gbps with a default of 2Gbps for the primary path and 1.5Gbps for the secondary path. The capability to set a path limit for the IOM1/2 can be seen when comparing Output 3 with Output 4; in the latter the “AvailBW” and “TotalBw” for the “PATH” shows the path limit.

In addition to setting the path limits in the bandwidth-policy, they can also be overridden on a given MDA.

```
config# card slot-number mda mda-slot
    ingress
        mcast-path-management
            primary-override
                path-limit megabits-per-second
            secondary-override
                path-limit megabits-per-second
```

The achievable capacity will be the minimum of the path’s path-limit and the plane’s per-mcast-plane-limit.

Ancillary Path

The ancillary path

The ancillary path allows managed multicast traffic to be forwarded through the switch fabric as unicast and so is not constrained to the path or plane capacities. This is achieved using ingress replication, in order to send a channel to multiple destination forwarding complexes (DFCs), the ingress forwarding complex creates and forwards one copy of each packet to each DFC connected to the switch fabric.

However, the total replication capacity available for the ancillary path is constrained to 5G to prevent it impacting the unicast or primary/secondary path capacities. This means that the total amount of ancillary capacity usable can be calculated from (note that the first copy sent is not included in this capacity, hence the “-1”):

5Gbps/(number_of_switch_fabric_DFCs – 1)

Taking an example shown later, if some channels ingress an IOM2 and egress 2 IOM3-XP (1 DFC each) and 1 IMM8 (1 DFC) to give a total of 3 egress DFCs, then total ancillary capacity available is

5Gbps/(3-1) = 2.5Gbps.

This would allow, for example, approximately 250 channels at 10Mbps each to use the ancillary path.

Due to the relationship between ancillary capacity and number of DFCs, the system will prefer the ancillary path as default whenever a channel enters an IOM1/2 and egresses on up to 3 DFCs. If there are 4 or more egress DFCs for the channel, then the primary path is preferred. The determination is performed on a per channel basis.

The configuration parameters relating to the primary and secondary paths are also available for the ancillary path.

```
config# mcast-management bandwidth-policy policy-name create
    ancillary-path
        queue-parameters
            cbs percentage
            hi-priority-only percent-of-mbs
            mbs percentage

config# mcast-management bandwidth-policy policy-name create
    ancillary-path
        path-limit megabits-per-second

config# card slot-number mda mda-slot
    ingress
        mcast-path-management
            ancillary-override
                path-limit megabits-per-second
```

IOM1/2 Planes

The capacity per plane for managed traffic is by default 2Gbps for a primary path and 1.8Gbps for a secondary path. Note that the default IOM1/2 secondary path limit is 1.5Gbps. A maximum of 2Gbps should be configured for either path type using the per-mcast-plane-limit (as shown for the IOM3-XP/IMM) when an IOM1/2 is being used with IMPM.

As the ancillary path does not use the switch fabric planes, there is no associated plane limit.

IOM1/2 Path to Plane Mapping

The hi-bw-mcast-src command function is the same for IOM1/2 line cards as for IOM3-XP/IMM line cards, as described above.

IMPM Operation on IOM1/2

This is exactly the same as the operation as for the IOM3-XP/IMM, see above.

IMPM Not Enabled

When IMPM is not enabled most multipoint traffic on an IOM1/2 and IOM3-XP/IMM can use only one primary path and one secondary path per forwarding complex. When ingress multipoint arrives it is placed on a multipoint queue and these queues are connected either to a primary path (if the queue is expedited) or a secondary path (if the queue is best-effort) depending on the ingress QOS classification applied. Standard ingress forwarding class/scheduling prioritization is used.

The capacity of the primary and secondary paths is 2Gbps, unless the system is a 7750 SR-12 or 7450 ESS-12 populated with SF/CPM4(s) and 100G FP IMMs in which case the capacity is 4Gbps.

In Output 1 and Output 2, the primary path is associated with the left-most plane and the secondary path is associated with the right-most plane for each line card.

There are exceptions to the above on the IOM3-XP/IMM line cards for

- Point-to-multipoint LSP IP multicast traffic
- Policed ingress routed IP multicast or VPLS broadcast, unknown or multicast traffic

Point-to-Multipoint (P2MP) LSP IP Multicast Traffic

IMPM will manage traffic on a P2MP LSP for any IP multicast channel that is delivered locally, for example, the system is a bud LSR. However, non-managed P2MP LSP IP multicast traffic will also make use of the primary paths, regardless of whether IMPM is enabled or not.

For each primary queue created in the MC IMPM Path pool, an additional queue is created to carry non-managed P2MP LSP IP multicast traffic. The non-managed P2MP LSP IP multicast traffic is automatically distributed across all primary paths based on a modulo N function of the 10 least significant bits of channel destination group address, where N is the number of primary paths. Note that the number of primary paths can be changed with IMPM enabled or disabled by applying a bandwidth-policy which sets the number of secondary paths.

Policed Ingress Routed IP Multicast or VPLS Broadcast, Unknown Or Multicast Traffic

Routed IP multicast or VPLS broadcast, unknown or multicast traffic passing through ingress hardware policers on the IOM3-XP/IMM can also use the IMPM managed queues, with IMPM enabled or disabled. If this traffic is best-effort (forwarding classes BE, L2, AF, L1) it will use the secondary paths, if it is expedited (forwarding classes H2, EF, H1, NC) it will use the primary paths. Note that this traffic uses the shared ingress policer-output-queues which have a fixed forwarding class to queue mapping).

When IMPM is not enabled, this traffic is non-managed and 1 secondary path plus 15 primary paths are available (the default). Consequently, extra capacity is only available for the expedited traffic, which could use up to 15 planes worth of switch fabric capacity. If extra capacity is required for best-effort traffic, a bandwidth-policy allocating more secondary paths can applied to the line card even without IMPM being enabled.

The policed ingress routed IP or VPLS broadcast, unknown or multicast traffic is distributed across the paths using a standard LAG hash algorithm (as described in the LAG and ECMP Hashing section in the Interface Configuration guide).

For both of these exceptions, it is recommended to reduce the managed traffic primary/secondary plane limits (using per-mcast-plane-limit) in order to allow for the non-managed traffic.

Show Output

This section includes the show output related to IMPM. The first part covers generic output and uses IOM3-XP/IMMs and chassis mode **d**. The second part includes an IOM2 and so uses chassis mode **a**.

IOM3-XP/IMM and Generic Output

The system has an IOM3-XP in slots 6 and 8, with an IMM8 is slot 7.

```
A:PE-1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned      Equipped         Admin   Operational        Comments
          Card-type        Card-type        State   State
-------------------------------------------------------------------------------
```

```
6         iom3-xp          iom3-xp          up      up
7         imm8-10gb-xfp    imm8-10gb-xfp    up      up
8         iom3-xp          iom3-xp          up      up
A         sfm4-12          sfm4-12          up      up/active
B         sfm4-12          sfm4-12          up      up/standby
===============================================================================
A:PE-1#
```

The status of IMPM on a given card can be shown as follows:

```
*A:PE-1# show card 6 detail
===============================================================================
Card 6
===============================================================================
Slot      Provisioned      Equipped         Admin   Operational        Comments
          Card-type        Card-type        State   State
-------------------------------------------------------------------------------
6         iom3-xp          iom3-xp          up      up

FP 1 Specific Data
    hi-bw-mc-srcEgress Alarm      : 2
    hi-bw-mc-srcEgress Group      : 0
    mc-path-mgmt Admin State      : In Service
    Ingress Bandwidth Policy      : default
```

IMPM is enabled on the fp, it is using the default bandwidth-policy and is using the default hi-bw-mcast-src group (0).

The MC Path Mgmt pool is created by default with the default settings.

```
*A:PE-1# show pools 6/1
===============================================================================

===============================================================================
Type    Id        App.    Pool Name                  Actual ResvCBS  PoolSize
                                                     Admin ResvCBS
-------------------------------------------------------------------------------
MDA     6/1       Acc-Ing MC Path Mgmt               18816           37632
                                                     50%
===============================================================================
*A:PE-1#
```

The default bandwidth-policy can be shown, giving the default parameters for the MC Path Pool and the associated queues.

```
*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
Limit              : 2000 mbps          Cbs                : 5.00
Mbs                : 7.00               High Priority      : 10
Secondary
Limit              : 1500 mbps          Cbs                : 30.00
Mbs                : 40.00              High Priority      : 10
Ancillary
Limit              : 5000 mbps          Cbs                : 65.00
Mbs                : 80.00              High Priority      : 10
T2-Primary
Cbs                : 5.00               Mbs                : 7.00
High Priority      : 10
T2-Secondary
Cbs                : 30.00              Mbs                : 40.00
High Priority      : 10                 Paths(Single/Dual) : 1/1
===============================================================================
Bandwidth Policies : 1
===============================================================================
*A:PE-1#
```

The defaults for the multicast-info-policy can be seen in configuration mode.

```
*A:PE-1# configure mcast-management
*A:PE-1>config>mcast-mgmt# info detail
----------------------------------------------
            multicast-info-policy "default" create
                no description
                bundle "default" create
                    no cong-priority-threshold
                    no description
                    no ecmp-opt-threshold
                    no admin-bw
                    no preference
                    no keepalive-override
                    no explicit-sf-path
                    bw-activity dynamic falling-delay 30
                    no primary-tunnel-interface
                exit
            exit
```

The paths/planes on an IOM3-XP/IMM can be shown here for card 6.

```
*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[6][0]: 0xf33b0a00
PATH:                                       PLANE:
Type SGs     InUseBW   AvailBW   TotalBw  ID   SGs   InUseBW   AvailBW   TotalBw
P      1           0         -         -   4     1         0   2000000   2000000
P      1           0         -         -   3     1         0   2000000   2000000
```

```
P      1           0         -         -   5     1         0   2000000   2000000
P      1           0         -         -   6     1         0   2000000   2000000
P      1           0         -         -   7     1         0   2000000   2000000
P      1           0         -         -   8     1         0   2000000   2000000
P      1           0         -         -   9     1         0   2000000   2000000
P      1           0         -         -  10     1         0   2000000   2000000
P      1           0         -         -  11     1         0   2000000   2000000
P      1           0         -         -  12     1         0   2000000   2000000
P      1           0         -         -  13     1         0   2000000   2000000
P      1           0         -         -  14     1         0   2000000   2000000
P      1           0         -         -  15     1         0   2000000   2000000
P      1           0         -         -  16     1         0   2000000   2000000
P      1           0         -         -   0     1         0   2000000   2000000
S      1           0         -         -   1     1         0   1800000   1800000
B      0           0         -         -   -     -         -         -         -
*A:PE-1#
```

Notice the plane total bandwidth is by default 2000Mbps for the primary paths and 1800Mbps for the secondary path, as can also be seen using this output.

```
*A:PE-1# show mcast-management chassis
===============================================================================
Chassis Information
===============================================================================
BW per MC plane               Single SFM  Dual SFM
-------------------------------------------------------------------------------
Primary Path                  2000        2000
Secondary Path                1800        1800
-------------------------------------------------------------------------------
MMRP Admin Mode               Disabled
MMRP Oper Mode                Disabled
Round Robin Inactive Records  Disabled
===============================================================================
*A:PE-1#
```

The Round Robin Inactive Records is disabled. The MMRP (Multiple MAC Registration Protocol) modes relate to the use of the MC Path Mgmt queues for MMRP traffic. When this is enabled, normal IMPM behavior is suspended so it is not in the scope of this configuration note.

A single channel (239.255.0.2) is now sent into interface int-IOM3-1 on port 6/2/1 with static IGMP joins on interfaces int-IMM8, int-IOM3-1 and int-IOM3-2. The current forwarding rate can be seen.

```
*A:PE-1# show router pim group detail
===============================================================================
PIM Source Group ipv4
===============================================================================
Group Address      : 239.255.0.2
Source Address     : 172.16.2.1
RP Address         : 192.0.2.1
Flags              : spt, rpt-prn-des   Type              : (S,G)
MRIB Next Hop      : 172.16.2.1
```

```
MRIB Src Flags     : direct             Keepalive Timer Exp: 0d 00:03:14
Up Time            : 0d 00:00:16        Resolved By        : rtable-u

Up JP State        : Joined             Up JP Expiry       : 0d 00:00:00
Up JP Rpt          : Pruned             Up JP Rpt Override : 0d 00:00:00

Register State     : Pruned             Register Stop Exp  : 0d 00:00:59
Reg From Anycast RP: No

Rpf Neighbor       : 172.16.2.1
Incoming Intf      : int-IOM3-1
Outgoing Intf List : int-IMM8, int-IOM3-1, int-IOM3-2

Curr Fwding Rate   : 9873.0 kbps
Forwarded Packets  : 18017              Discarded Packets  : 0
Forwarded Octets   : 24899494           RPF Mismatches     : 0
Spt threshold      : 0 kbps             ECMP opt threshold : 7
Admin bandwidth    : 1 kbps

===============================================================================
```

From the two sets of output below it can be seen that this is using the default primary path and switch fabric plane 4.

```
*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[6][0]: 0xf33b0a00
PATH:                                       PLANE:
Type SGs     InUseBW   AvailBW   TotalBw  ID   SGs   InUseBW   AvailBW   TotalBw
P      2        9895         -         -   4     2      9895   1990105   2000000
P      1           0         -         -   3     1         0   2000000   2000000
P      1           0         -         -   5     1         0   2000000   2000000
P      1           0         -         -   6     1         0   2000000   2000000
P      1           0         -         -   7     1         0   2000000   2000000
P      1           0         -         -   8     1         0   2000000   2000000
P      1           0         -         -   9     1         0   2000000   2000000
P      1           0         -         -  10     1         0   2000000   2000000
P      1           0         -         -  11     1         0   2000000   2000000
P      1           0         -         -  12     1         0   2000000   2000000
P      1           0         -         -  13     1         0   2000000   2000000
P      1           0         -         -  14     1         0   2000000   2000000
P      1           0         -         -  15     1         0   2000000   2000000
P      1           0         -         -  16     1         0   2000000   2000000
P      1           0         -         -   0     1         0   2000000   2000000
S      1           0         -         -   1     1         0   1800000   1800000
B      0           0         -         -   -     -         -         -         -
*A:PE-1#

*A:PE-1# show system switch-fabric high-bandwidth-multicast
===============================================================================
Switch Fabric
===============================================================================
         Cap:               Planes:
Slot/Mda Min  Max  Hbm Grp  Hi | Lo
-------------------------------------------------------------------------------
 6/1     100% 100% No  0     4  3  5  6  7  8  9 10 11 12 13 14 15 16  0 |  1
 7/1     100% 100% No  0    19 17 20 21 22 23 24 25 26 27 28 29 30 31 32 | 33
 8/1     100% 100% No  0    35 34 36 37 38 39 40 41 42 43 44 45 46 47  0 |  1
```

```
 A       100% 100% No  0     2 |  2
 B       100% 100% No  0     2 |  2
===============================================================================
*A:PE-1#
```

The information about the channel can be seen using this command.

```
*A:PE-1# show mcast-management
    channel [router router-instance|vpls service-id|service-name service-name]
            [mda slot[/mda]]
            [group ip-address [source ip-address]]
            [path path-type]
            [detail]
```

The output for the channel being sent is as follows.

```
*A:PE-1# show mcast-management channel
===============================================================================
Multicast Channels
===============================================================================
Legend :  D - Dynamic  E - Explicit
===============================================================================
Source Address                              Slot/Cpx  Current-Bw  Path      D/E
 Group Address                                         Highest-Bw  Plane
-------------------------------------------------------------------------------
172.16.2.1                                  6/1       9873        Primary    D
 239.255.0.2                                           9873        4
===============================================================================
Multicast Channels : 1
===============================================================================
*A:PE-1#
```

```
*A:PE-1# show mcast-management channel detail
===============================================================================
Multicast Channels
===============================================================================
-------------------------------------------------------------------------------
Source Address     : 172.16.2.1
Group Address      : 239.255.0.2
-------------------------------------------------------------------------------
Slot/Complex       : 6/1                Current Bw         : 9873 kbps
Dynamic/Explicit   : Dynamic            Current Path       : Primary
Oper Admin Bw      : 0 kbps             Current Plane      : 4
Ing last highest   : 9873               Preference         : 0
Black-hole rate    : None               Ing sec highest    : 9873
Time remaining     : 30 seconds         Blackhole          : No
===============================================================================
Multicast Channels : 1
===============================================================================
*A:PE-1#
```

The channel is using the dynamic bandwidth activity measurement and the current bandwidth, last highest and second last highest rates are shown (which are the same as this traffic is from a traffic generator).

The Time remaining is the time remaining in the current falling-delay period. This is reset to the falling-delay every time the last highest bandwidth gets updated, when it reaches zero the value of last highest bandwidth will be replaced with second highest bandwidth and the second highest bandwidth will be set to the value of current bandwidth.

The Oper Admin Bw displays the value used for the admin-bw for this channel.

A subset of this information can be seen using this tools command.

```
*A:PE-1# tools dump mcast-path-mgr channels slot 6
===============================================================================
Slot: 6 Complex: 0
===============================================================================
Source address                          CurrBw   Plane PathType  Path Pref
 Group address                          PathBw   Repl  Exp       BlkHoleBw
-------------------------------------------------------------------------------
172.16.2.1                              9873     4     primary   0    0
 239.255.0.2                            9873     2     none      0
Unmanaged traffic                       0        4     primary   0    8
 slot: 6 cmplx: 0 path: 0               0        0     none      0
Unmanaged traffic                       0        3     primary   1    8
 slot: 6 cmplx: 0 path: 1               0        0     none      0
Unmanaged traffic                       0        5     primary   2    8
 slot: 6 cmplx: 0 path: 2               0        0     none      0
Unmanaged traffic                       0        6     primary   3    8
 slot: 6 cmplx: 0 path: 3               0        0     none      0
Unmanaged traffic                       0        7     primary   4    8
 slot: 6 cmplx: 0 path: 4               0        0     none      0
Unmanaged traffic                       0        8     primary   5    8
 slot: 6 cmplx: 0 path: 5               0        0     none      0
Unmanaged traffic                       0        9     primary   6    8
 slot: 6 cmplx: 0 path: 6               0        0     none      0
Unmanaged traffic                       0        10    primary   7    8
 slot: 6 cmplx: 0 path: 7               0        0     none      0
Unmanaged traffic                       0        11    primary   8    8
 slot: 6 cmplx: 0 path: 8               0        0     none      0
Unmanaged traffic                       0        12    primary   9    8
 slot: 6 cmplx: 0 path: 9               0        0     none      0
Unmanaged traffic                       0        13    primary   10   8
 slot: 6 cmplx: 0 path: 10              0        0     none      0
Unmanaged traffic                       0        14    primary   11   8
 slot: 6 cmplx: 0 path: 11              0        0     none      0
Unmanaged traffic                       0        15    primary   12   8
 slot: 6 cmplx: 0 path: 12              0        0     none      0
Unmanaged traffic                       0        16    primary   13   8
 slot: 6 cmplx: 0 path: 13              0        0     none      0
Unmanaged traffic                       0        0     primary   14   8
 slot: 6 cmplx: 0 path: 14              0        0     none      0
Unmanaged traffic                       0        1     secondary 15   8
 slot: 6 cmplx: 0 path: 15              0        0     none      0
*A:PE-1#
```

The bandwidth activity monitoring is now changed to use an admin-bw of 12Mbps with a blackhole rate of 15Mbps.

```
*A:PE-1# configure mcast-management
*A:PE-1>config>mcast-mgmt# info
----------------------------------------------
            bandwidth-policy "bandwidth-policy-1" create
                admin-bw-threshold 8000
            exit
            multicast-info-policy "multicast-info-policy-1" create
                bundle "default" create
                exit
                bundle "bundle-1" create
                    channel "239.255.0.1" "239.255.0.16" create
                        admin-bw 12000
                        bw-activity use-admin-bw black-hole-rate 15000
                    exit
                exit
            exit
----------------------------------------------
*A:PE-1>config>mcast-mgmt# exit all


*A:PE-1# show mcast-management channel group 239.255.0.2 detail
===============================================================================
Multicast Channels
===============================================================================
-------------------------------------------------------------------------------
Source Address     : 172.16.2.1
Group Address      : 239.255.0.2
-------------------------------------------------------------------------------
Slot/Complex       : 6/1                Current Bw         : 9873 kbps
Dynamic/Explicit   : Dynamic            Current Path       : Primary
Oper Admin Bw      : 12000 kbps         Current Plane      : 4
Ing last highest   : 12000              Preference         : 0
Black-hole rate    : 15000 kbps         Ing sec highest    : 12000
Time remaining     : 30 seconds         Blackhole          : No
===============================================================================
Multicast Channels : 1
===============================================================================
*A:PE-1#


*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[6][0]: 0xf33b0a00
PATH:                                       PLANE:
Type SGs     InUseBW   AvailBW   TotalBw  ID   SGs   InUseBW   AvailBW   TotalBw
P      2       12000         -         -   4     2     12000   1988000   2000000
P      1           0         -         -   3     1         0   2000000   2000000
P      1           0         -         -   5     1         0   2000000   2000000
P      1           0         -         -   6     1         0   2000000   2000000
P      1           0         -         -   7     1         0   2000000   2000000
P      1           0         -         -   8     1         0   2000000   2000000
P      1           0         -         -   9     1         0   2000000   2000000
P      1           0         -         -  10     1         0   2000000   2000000
P      1           0         -         -  11     1         0   2000000   2000000
P      1           0         -         -  12     1         0   2000000   2000000
P      1           0         -         -  13     1         0   2000000   2000000
```

```
P      1           0         -         -  14     1         0   2000000   2000000
P      1           0         -         -  15     1         0   2000000   2000000
P      1           0         -         -  16     1         0   2000000   2000000
P      1           0         -         -   0     1         0   2000000   2000000
S      1           0         -         -   1     1         0   1800000   1800000
B      0           0         -         -   -     -         -         -         -
*A:PE-1#
```

Now the system treats the channel as though it is using 12Mbps capacity even though its current rate has not changed.

If the rate is increased above the blackhole rate, the channel is blackholed and an alarm is generated.

```
*A:PE-1#
11 2011/10/21 01:40:13.21 UTC MINOR: MCPATH #2001 Base Black-hole-rate is reached
"Channel (172.16.2.1,239.255.0.2) for vRtr instance 1 slot/cplx 6/
1 has been blackholed."
*A:PE-1# show mcast-management channel group 239.255.0.2 detail
===============================================================================
Multicast Channels
===============================================================================
-------------------------------------------------------------------------------
Source Address     : 172.16.2.1
Group Address      : 239.255.0.2
-------------------------------------------------------------------------------
Slot/Complex       : 6/1                Current Bw         : 19458 kbps
Dynamic/Explicit   : Dynamic            Current Path       : Blackhole
Oper Admin Bw      : 12000 kbps         Current Plane      : N/A
Ing last highest   : 19480              Preference         : 0
Black-hole rate    : 15000 kbps         Ing sec highest    : 19469
Time remaining     : 23 seconds         Blackhole          : Yes
===============================================================================
Multicast Channels : 1
===============================================================================
*A:PE-1#
*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[6][0]: 0xf33b0a00
PATH:                                    PLANE:
Type SGs     InUseBW   AvailBW   TotalBw  ID   SGs   InUseBW   AvailBW   TotalBw
P      1           0         -         -   4     1         0   2000000   2000000
P      1           0         -         -   3     1         0   2000000   2000000
P      1           0         -         -   5     1         0   2000000   2000000
P      1           0         -         -   6     1         0   2000000   2000000
P      1           0         -         -   7     1         0   2000000   2000000
P      1           0         -         -   8     1         0   2000000   2000000
P      1           0         -         -   9     1         0   2000000   2000000
P      1           0         -         -  10     1         0   2000000   2000000
P      1           0         -         -  11     1         0   2000000   2000000
P      1           0         -         -  12     1         0   2000000   2000000
P      1           0         -         -  13     1         0   2000000   2000000
P      1           0         -         -  14     1         0   2000000   2000000
P      1           0         -         -  15     1         0   2000000   2000000
P      1           0         -         -  16     1         0   2000000   2000000
P      1           0         -         -   0     1         0   2000000   2000000
S      1           0         -         -   1     1         0   1800000   1800000
```

```
B      1       19480         -         -   -     -         -         -         -
*A:PE-1#
```

The output displayed above shows an alarm generated for a channel being blackholed due to the channel rate reaching the configured black-hole-rate. The example output below is an alarm for a channel being blackholed due to insufficient bandwidth being available to it.

```
7 2011/10/22 21:53:43.54 UTC MINOR: MCPATH #2001 Base No bandwidth available
"Channel (172.16.2.1,239.255.0.2) for vRtr instance 1 slot/cplx 6/
1 has been blackholed."
```

Note that the following alarm relates to a dummy channel used to account for the unmanaged traffic. However, this traffic is never actually blackholed.

```
6 2011/10/21 00:27:58.00 UTC MINOR: MCPATH #2002 Base
"Channel (0.0.0.0,0.6.0.0) for unknown value (2) instance 0 slot/cplx 6/
1 is no longer being blackholed."
```

Alarms are also generated when all paths of a given type reach certain thresholds.

For primary and secondary paths, the two path range limit thresholds are

- Full: less than 5% capacity is available

```
9 2011/10/24 22:53:27.02 UTC MINOR: MCPATH #2003 Base
"The available bandwidth on secondary path on slot/cplx 6/
1 has reached its maximum limit."
```

- Not full: more than 10% of the path capacity is available.

```
10 2011/10/24 22:53:48.02 UTC MINOR: MCPATH #2004 Base
"The available bandwidth on secondary path on slot/cplx 6/1 is within range limits."
```

A maximum of one alarm is generated for each event (blackhole start/stop, path full/not full) within a 3 second period. So, for example, if multiple channels are blackholed within the same 3 second period only one alarm will be generated (for the first event).

The effect of using the hi-bw-mcast-src command is illustrated below. Firstly, line card 6 is configured into group 1.

```
*A:PE-1# configure card 6 fp hi-bw-mcast-src group 1 alarm
*A:PE-1# show system switch-fabric high-bandwidth-multicast
===============================================================================
Switch Fabric
===============================================================================
         Cap:               Planes:
Slot/Mda Min  Max  Hbm Grp  Hi | Lo
-------------------------------------------------------------------------------
 6/1     100% 100% Yes 1     4  3  5  6  7  8  9 10 11 12 13 14 15 16   0 |  1
```

```
 7/1     100% 100% No  0    19 17 20 21 22 23 24 25 26 27 28 29 30 31 32 | 33
 8/1     100% 100% No  0    35 34 36 37 38 39 40 41 42 43 44 45 46 47  0 |  1
 A       100% 100% No  0     2 |  2
 B       100% 100% No  0     2 |  2
===============================================================================
*A:PE-1#
```

The plane assignment has not changed (though it is possible that the system could re-arrange the planes used by card 6) and there are still planes (0,1) shared between card 6 and card 8.

Now card 8 is configured into group 2 (note that IMPM is only enabled on card 6 here).

```
*A:PE-1# configure card 8 fp hi-bw-mcast-src group 2 alarm
*A:PE-1# show system switch-fabric high-bandwidth-multicast
===============================================================================
Switch Fabric
===============================================================================
         Cap:               Planes:
Slot/Mda Min  Max  Hbm Grp  Hi | Lo
-------------------------------------------------------------------------------
 6/1     100% 100% Yes 1     4  3  5  6  7  8  9 10 11 12 13 14 15 16  0 |  1
 7/1     100% 100% No  0    19 17 20 21 22 23 24 25 26 27 28 29 30 31 32 | 33
 8/1     100% 100% Yes 2    35 34 36 37 38 39 40 41 42 43 44 45 46 47 17 | 19
 A       100% 100% No  0     2 |  2
 B       100% 100% No  0     2 |  2
===============================================================================
*A:PE-1#
```

There are no longer planes shared between cards 6 and 8.

If card 7 is configured into group 3, the following is seen.

```
*A:PE-1# configure card 7 fp hi-bw-mcast-src group 3 alarm
*A:PE-1#
7 2011/10/21 00:35:50.95 UTC MINOR: CHASSIS #2052 Base Mda 6/1
"Class MDA Module : Plane shared by multiple multicast high bandwidth taps"

8 2011/10/21 00:35:50.95 UTC MINOR: CHASSIS #2052 Base Mda 6/2
"Class MDA Module : Plane shared by multiple multicast high bandwidth taps"

9 2011/10/21 00:35:50.97 UTC MINOR: CHASSIS #2052 Base Mda 7/1
"Class MDA Module : Plane shared by multiple multicast high bandwidth taps"

10 2011/10/21 00:35:50.97 UTC MINOR: CHASSIS #2052 Base Mda 7/2
"Class MDA Module : Plane shared by multiple multicast high bandwidth taps"


*A:PE-1# show system switch-fabric high-bandwidth-multicast
===============================================================================
Switch Fabric
===============================================================================
         Cap:               Planes:
Slot/Mda Min  Max  Hbm Grp  Hi | Lo
```

```
-------------------------------------------------------------------------------
 6/1     100% 100% Yes 1     4  3  5  6  7  8  9 10 11 12 13 14 15 16  0 |  1
 7/1     100% 100% Yes 3    21 20 22 23 24 25 26 27 28 29 30 31 32 33  0 |  1
 8/1     100% 100% Yes 2    35 34 36 37 38 39 40 41 42 43 44 45 46 47 17 | 19
 A       100% 100% No  0     2 |  2
 B       100% 100% No  0     2 |  2
===============================================================================
*A:PE-1#
```

There are insufficient planes to allow each card/group to have dedicated planes. Planes 0 and 1 are still shared between cards 6 and 7, generating the associated alarms.

A common example of the use of the **hi-bw-mcast-src** command would be when cards 6 and 8 have uplink ports on which high bandwidth multicast channels could be received. It would be desired to have these cards use different planes. To achieve this, card 7 could be configured into group 1, as follows.

```
*A:PE-1# configure card 7 fp hi-bw-mcast-src group 1 alarm
*A:PE-1# show system switch-fabric high-bandwidth-multicast
===============================================================================
Switch Fabric
===============================================================================
         Cap:              Planes:
Slot/Mda Min  Max  Hbm Grp  Hi | Lo
-------------------------------------------------------------------------------
 6/1     100% 100% Yes 1     4  3  5  6  7  8  9 10 11 12 13 14 15 16  0 |  1
 7/1     100% 100% Yes 1     4  3  5  6  7  8  9 10 11 12 13 14 15 16  0 |  1
 8/1     100% 100% Yes 2    35 34 36 37 38 39 40 41 42 43 44 45 46 47 17 | 19
 A       100% 100% No  0     2 |  2
 B       100% 100% No  0     2 |  2
===============================================================================
*A:PE-1#
```

Now it can be seen that card 7 shares the same planes as card 6, but more importantly card 6 has no planes in common with card 8.

Note that when traffic is received on card 6, it will also be seen on the same plane (not path) on card 7. In the example below, traffic can be seen on plane 4 which is used by both cards 6 and 7, but only card 6 has non-zero InUseBW path capacity.

```
*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[6][0]: 0xf33b0a00
PATH:                                       PLANE:
Type SGs     InUseBW   AvailBW   TotalBw  ID   SGs   InUseBW   AvailBW   TotalBw
P      2        9707         -         -   4     3      9707   1990293   2000000
P      1           0         -         -   3     2         0   2000000   2000000
...
McPathMgr[7][0]: 0xf33b3198
PATH:                                       PLANE:
Type SGs     InUseBW   AvailBW   TotalBw  ID   SGs   InUseBW   AvailBW   TotalBw
P      1           0         -         -   4     3      9707   1990293   2000000
```

```
P      1          0         -         -   3     2         0   2000000   2000000
```

When IMPM managed traffic is received on SAPs (in an IES, VPLS or VPRN service) it can be seen against a specific queue counter (Off. Managed) in the SAP stats. The following output shows where sap 7/2/1:3 belongs to a VPLS service using igmp-snooping. A similar counter is not available for policer statistics.

```
*A:PE-1# show service id 2 sap 7/2/1:3 stats
===============================================================================
Service Access Points(SAP)
===============================================================================
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LoPrio           : 0                       0
Dro. HiPrio           : 0                       0
Dro. LoPrio           : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio           : 0                       0
Off. LoPrio           : 0                       0
Off. Managed          : 149410                  209771640
Dro. HiPrio           : 0                       0
Dro. LoPrio           : 0                       0
For. InProf           : 149410                  209771640
For. OutProf          : 0                       0

Egress Queue 1
For. InProf           : 0                       0
For. OutProf          : 0                       0
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
===============================================================================
*A:PE-1#
```

IOM1/2 Specific Output

The system is configured with the following cards and is in chassis mode **a**. As can be seen, an IOM2 is in slot 5.

```
A:PE-1# show card
===============================================================================
Card Summary
===============================================================================
Slot      Provisioned      Equipped         Admin   Operational        Comments
          Card-type        Card-type        State   State
```

```
-------------------------------------------------------------------------------
5         iom2-20g         iom2-20g         up      up
6         iom3-xp          iom3-xp          up      up
7         imm8-10gb-xfp    imm8-10gb-xfp    up      up
8         iom3-xp          iom3-xp          up      up
A         sfm4-12          sfm4-12          up      up/active
B         sfm4-12          sfm4-12          up      up/standby
===============================================================================
A:PE-1#
```

IMPM is enabled on MDA 1 and 2 of the IOM2 in slot 5, with a primary, secondary and ancillary path.

```
A:PE-1# show mcast-management mda
===============================================================================
MDA Summary
===============================================================================
S/C  Policy                           Type            In-use-Bw          Admin
-------------------------------------------------------------------------------
5/1  default                          Primary         0 Kbps             up
     default                          Secondary       0 Kbps             up
     default                          Ancillary       0 Kbps             up
5/2  default                          Primary         0 Kbps             up
     default                          Secondary       0 Kbps             up
     default                          Ancillary       0 Kbps             up
6/2  default                          Primary         0 Kbps             down
     default                          Secondary       0 Kbps             down
     default                          Ancillary       0 Kbps             down
7/1  default                          Primary         0 Kbps             down
     default                          Secondary       0 Kbps             down
     default                          Ancillary       0 Kbps             down
7/2  default                          Primary         0 Kbps             down
     default                          Secondary       0 Kbps             down
     default                          Ancillary       0 Kbps             down
8/2  default                          Primary         0 Kbps             down
     default                          Secondary       0 Kbps             down
     default                          Ancillary       0 Kbps             down
===============================================================================
A:PE-1#
```

The path/plane usage can be shown.

```
*A:PE-1# show system switch-fabric high-bandwidth-multicast
===============================================================================
Switch Fabric
===============================================================================
         Cap:               Planes:
Slot/Mda Min  Max  Hbm Grp  Hi | Lo
-------------------------------------------------------------------------------
 5/1     100% 100% No  0     1 |  0
 5/2     100% 100% No  0     4 |  3
 6/1     100% 100% No  0     6  5  7  8  9 10 11 12 13 14 15  0  1  3  4 |  5
 7/1     100% 100% No  0     7  6  8  9 10 11 12 13 14 15  0  1  3  4  5 |  6
 8/1     100% 100% No  0     8  7  9 10 11 12 13 14 15  0  1  3  4  5  6 |  7
 A       100% 100% No  0     2 |  2
 B       100% 100% No  0     2 |  2
===============================================================================
```

```
*A:PE-1#


*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[5][0]: 0xf33b0a00
PATH:                                       PLANE:
Type SGs     InUseBW   AvailBW   TotalBw  ID   SGs   InUseBW   AvailBW   TotalBw
P      1           0   2000000   2000000   1     1         0   2000000   2000000
S      1           0   1500000   1500000   0     1         0   1800000   1800000
A      0           0   5000000   5000000   -     -         -         -         -
B      0           0         -         -   -     -         -         -         -
McPathMgr[5][1]: 0xf33b3198
PATH:                                       PLANE:
Type SGs     InUseBW   AvailBW   TotalBw  ID   SGs   InUseBW   AvailBW   TotalBw
P      1           0   2000000   2000000   4     1         0   2000000   2000000
S      1           0   1500000   1500000   3     1         0   1800000   1800000
A      0           0   5000000   5000000   -     -         -         -         -
B      0           0         -         -   -     -         -         -         -
*A:PE-1#
```

The path range limit alarm thresholds for the ancillary path are

- Full: less than 2% capacity is available
- Not full: more than 4% of the path capacity is available.

A single channel (239.255.0.1) is now sent into interface int-IOM2 on port 5/2/1 with static IGMP joins on interfaces int-IMM8, int-IOM3-1 and int-IOM3-2. The current forwarding rate can be seen.

```
*A:PE-1# show router pim group 239.255.0.1 detail
===============================================================================
PIM Source Group ipv4
===============================================================================
Group Address      : 239.255.0.1
Source Address     : 172.16.1.1
RP Address         : 192.0.2.1
Flags              : spt, rpt-prn-des   Type               : (S,G)
MRIB Next Hop      : 172.16.1.1
MRIB Src Flags     : direct             Keepalive Timer Exp: 0d 00:02:44
Up Time            : 0d 00:07:45        Resolved By        : rtable-u

Up JP State        : Joined             Up JP Expiry       : 0d 00:00:00
Up JP Rpt          : Pruned             Up JP Rpt Override : 0d 00:00:00

Register State     : Pruned             Register Stop Exp  : 0d 00:00:32
Reg From Anycast RP: No

Rpf Neighbor       : 172.16.1.1
Incoming Intf      : int-IOM2
Outgoing Intf List : int-IMM8, int-IOM3-1, int-IOM3-2

Curr Fwding Rate   : 9734.8 kbps
Forwarded Packets  : 591874             Discarded Packets  : 0
Forwarded Octets   : 817969868          RPF Mismatches     : 0
Spt threshold      : 0 kbps             ECMP opt threshold : 7
```

```
Admin bandwidth    : 1 kbps
-------------------------------------------------------------------------------
```

As there are only 3 (<4) DFCs, the ancillary path is used.

```
*A:PE-1# show mcast-management channel
===============================================================================
Multicast Channels
===============================================================================
Legend :  D - Dynamic  E - Explicit
===============================================================================
Source Address                              Slot/Cpx  Current-Bw  Path      D/E
 Group Address                                         Highest-Bw  Plane
-------------------------------------------------------------------------------
172.16.1.1                                  5/2       9729        Ancillary  D
 239.255.0.1                                          9740        -
===============================================================================
Multicast Channels : 1
===============================================================================
*A:PE-1#


*A:PE-1# show mcast-management channel detail
===============================================================================
Multicast Channels
===============================================================================
-------------------------------------------------------------------------------
Source Address     : 172.16.1.1
Group Address      : 239.255.0.1
-------------------------------------------------------------------------------
Slot/Complex       : 5/2               Current Bw         : 9729 kbps
Dynamic/Explicit   : Dynamic           Current Path       : Ancillary
Oper Admin Bw      : 0 kbps            Current Plane      : N/A
Ing last highest   : 9740              Preference         : 0
Black-hole rate    : None              Ing sec highest    : 9740
Time remaining     : 27 seconds        Blackhole          : No
===============================================================================
Multicast Channels : 1
===============================================================================
*A:PE-1#
```

If another join caused this traffic to be switched via an additional DFC, the system would place the channel on the primary path.

The ancillary path being used can also be seen as follows.

```
*A:PE-1# tools dump mcast-path-mgr cpm
McPathMgr[5][0]: 0xf33b0a00
PATH:                                       PLANE:
Type SGs     InUseBW   AvailBW   TotalBw  ID   SGs   InUseBW   AvailBW   TotalBw
P      1           0   2000000   2000000   1     1         0   2000000   2000000
S      1           0   1500000   1500000   0     1         0   1800000   1800000
A      0           0   5000000   5000000   -     -         -         -         -
B      0           0         -         -   -     -         -         -         -
McPathMgr[5][1]: 0xf33b3198
```

```
PATH:                                       PLANE:
Type SGs     InUseBW   AvailBW   TotalBw  ID   SGs   InUseBW   AvailBW   TotalBw
P      1           0   2000000   2000000   4     1         0   2000000   2000000
S      1           0   1500000   1500000   3     1         0   1800000   1800000
A      1       19480   4980520   5000000   -     -         -         -         -
B      0           0         -         -   -     -         -         -         -
*A:PE-1#
```

Note that the bandwidth shown on the ancillary path on the second MDA is approximately two times that of the ingress traffic, this matches the algorithm described earlier for the ancillary path. This can also be seen in the next output, there is the original channel traffic plus two replications (Repl).

```
*A:PE-1# tools dump mcast-path-mgr channels
===============================================================================
Slot: 5 Complex: 0
===============================================================================
Source address                          CurrBw   Plane PathType  Path Pref
 Group address                          PathBw   Repl  Exp       BlkHoleBw
-------------------------------------------------------------------------------
Unmanaged traffic                       0        1     primary   0    8
 slot: 5 cmplx: 0 path: 0               0        0     none      0
Unmanaged traffic                       0        0     secondary 1    8
 slot: 5 cmplx: 0 path: 1               0        0     none      0
===============================================================================
Slot: 5 Complex: 1
===============================================================================
Source address                          CurrBw   Plane PathType  Path Pref
 Group address                          PathBw   Repl  Exp       BlkHoleBw
-------------------------------------------------------------------------------
172.16.1.1                              9740     48    ancillary 16   0
 239.255.0.1                            19480    2     none      0
Unmanaged traffic                       0        4     primary   0    8
 slot: 5 cmplx: 1 path: 0               0        0     none      0
Unmanaged traffic                       0        3     secondary 1    8
 slot: 5 cmplx: 1 path: 1               0        0     none      0
*A:PE-1#
```

Conclusion

This chapter has described the configuration of Ingress Multicast Path Management which optimizes IPv4 and IPv6 multicast capacity to achieve the maximum system-wide IP multicast throughput. It can be used for both routed IPv4/IPv6 and VPLS (IGMP and PIM) snooped IPv4 multicast groups, which usually relate to the distribution of IP TV channels.

IPoE Sessions

Applicability

This chapter is applicable to the 7750 SR series and was tested on SR OS release 13.0.R7. Chassis-mode C or higher must be used.

IPoE sessions require a Routed CO environment with Enhanced Subscriber Management (ESM) enabled.

Overview

Definitions

Grouping a number of subscriber hosts with different IP stacks belonging to the same end user device in a single IPoE session simplifies operations; see Figure 154.

In this way, IPoE sessions provide behavior similar to PPPoE sessions for authentication, mid-session changes, and accounting.

Figure 154 **IPoE Session**

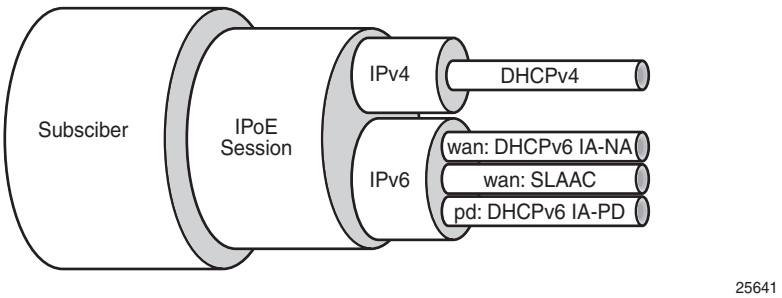

The hosts (IP stacks) associated with a single IPoE session share ESM data, such as the subscriber ID, the sub-profile, the SLA profile, and so on. The shared ESM data is fetched and cached when the first host for that session connects: only a single authentication is needed. This requires ESM to be enabled; see the ESM basics chapter for more information.

An IPoE session can have one IPoEv4 host, up to two IPoEv6 wan hosts (one DHCPv6 host and a SLAAC host), and one IPoEv6 pd host. Hosts with the same SAP, MAC address, and optionally the same circuit ID (CID) or remote ID (RID), are grouped in a single IPoE session, using that combination as a key to the IPoE session data.

Authentication occurs when the first host for that IPoE session is created. Subsequent hosts belonging to the same IPoE session do not require additional authentication. For instantiating these subsequent hosts, SR OS uses the cached ESM data that was fetched while authenticating and instantiating the first host.

Mid-session changes are typically triggered by RADIUS CoA or Diameter Gx RAR messages, and automatically apply to all hosts associated with the IPoE session. A re-authentication could also lead to policy changes, but the changes are triggered through host renewal messages.

As well as queue instance and host accounting, RADIUS session accounting can be enabled for IPoE sessions. An accounting session identity (ASID) is created when an IPoE session is started.

An IPoE session is created when the first host is created, and the IPoE session is deleted when the last host is deleted. IPoE session creation is always protocol triggered. IPoE session deletion is triggered through protocol (DHCPv4, DHCPv6 release, or expiration of a lease) or through policy (idle-timeout, session-timeout, clear command, and so on).

IPoE sessions require the Routed CO model, and are supported on regular as well as on capture and managed SAPs.

Trigger Packets

Unlike PPP, where PPP sessions have clear and unique triggers that start (PADS) and stop (PADT) a PPP session, IPoE does not have unique triggers that start and stop an IPoE session.

IPoE sessions are created when the following trigger packets on ESM-enabled SAPs are received:

- DHCPv4 discover/request
- DHCPv6 solicit/request (native)
- DHCPv6 solicit/request (single relay)
- DHCPv6 solicit/request (double relay)
- Router Solicitation

No IPoE sessions are created on reception of ARP requests or CoA messages.

IPoE Session Key

The key to the IPoE session data is a combination of the SAP ID, the MAC address, and optionally the CID or the RID, as defined by the ipoe-session-policy:

```
configure
    subscriber-mgmt
        ipoe-session-policy <pol-name>
            session-key sap mac [cid] [rid]
```

The ipoe-session-policy name *default* is reserved for future use.

The CID and RID are mutually exclusive; see Figure 155:

- The CID corresponds to DHCPv4 option 82, sub-option 1 (Circuit-ID) and to DHCPv6 option 18 (Interface-ID).
- The RID corresponds to DHCPv4 option 82, sub-option 2 (Remote-ID) and to DHCPv6 option 37 (the remote-id field of the Relay Agent Remote-ID, excluding the enterprise-number field).

Figure 155 **IPoE Session Key**

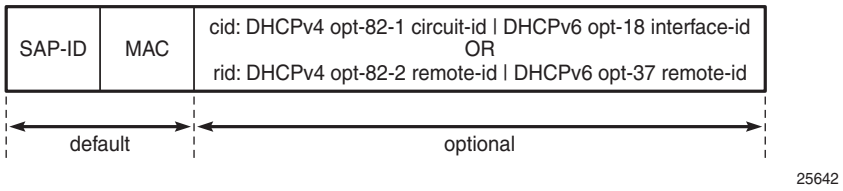

When an IPoE session trigger packet is received, the IPoE session key is validated, ensuring that no field is missing. For example, if the key requires the CID or RID, and a device connects without CID or RID, the IPoE session setup and the host setup fail. Therefore, the CID or RID should only be part of the key when all devices include this parameter in the trigger packets.

If no IPoE session exists for a session key derived from a trigger packet, an IPoE session is created. If an IPoE session exists, a new host is created and added to the existing IPoE session, on condition that the host type is compatible with the already associated host types.

IPoE Session Authentication

Authenticating IPoE sessions requires generic identification parameters, which must be supported in both IPv4 and IPv6, so some restrictions apply.

LUDB Authentication

When using an LUDB for IPoE session authentication, all of the host-identification criteria can be used, except for the following:

- Option 60 - DHCPv4 only

LUDB entries containing option 60 are ignored while scanning the LUDB for a matching entry.

AAA/RADIUS Authentication

When using AAA/RADIUS for IPoE session authentication, all username formats can be used, except for the following:

- dhcp-client-vendor-opts - DHCPv4 only
- mac-giaddr - DHCPv4 only
- ppp-user-name - PPP only

IPoE Session Creation

To ensure successful creation of an IPoE session, the following conditions must be met:

- the IPoE session key must be valid
- the session limits should not be exceeded
- an Accounting Session Identity (ASID) must be allocated
- the first ESM host must be successfully authenticated

Figure 156 shows the high-level flow of the IPoE session creation process. When a new trigger is received for an existing IPoE session, a new host is created and added to the session, on condition that the new host is compatible with the already associated hosts.

Figure 156 **IPoE Session Creation Flow**

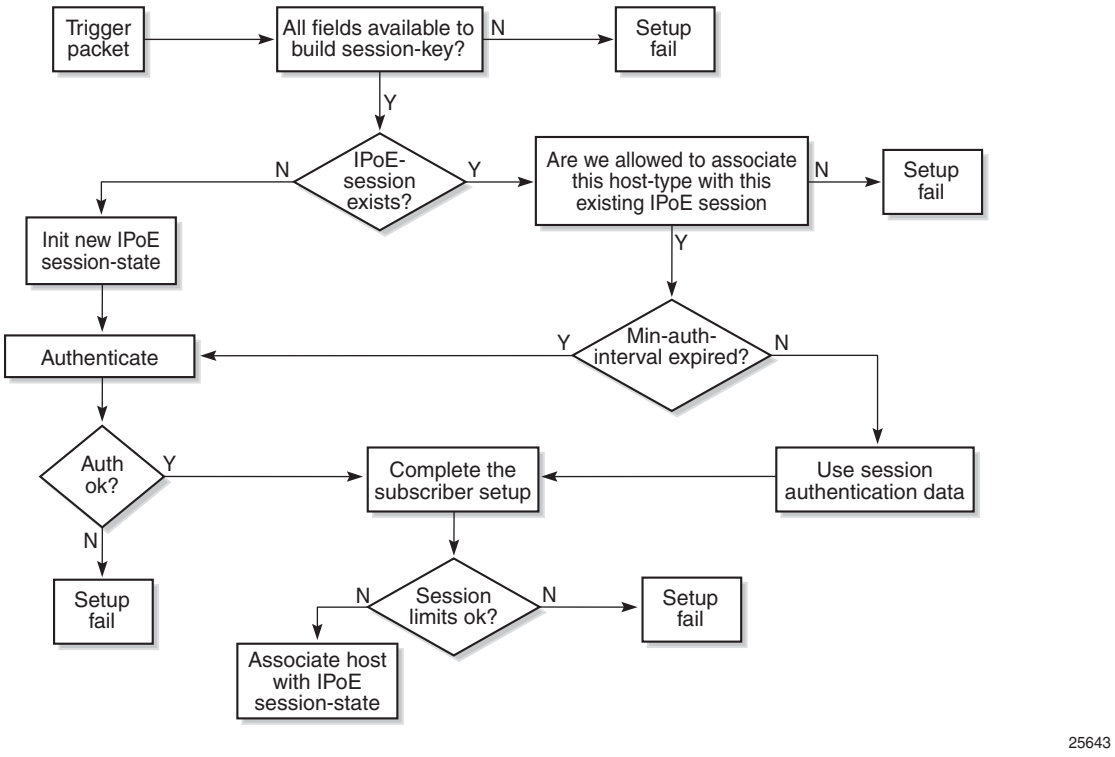

Figure 157 shows an example where a device initiates DHCPv6 first. The data fetched while authenticating this device through AAA/RADIUS is cached. When the same device (leading to the same key) later initiates DHCPv4, the cached IPoE session data can be used, so no further AAA/RADIUS access is needed. The sequence could also first initiate DHCPv4, then DHCPv6. The result would be the same.

Figure 157 **IPoE Session Creation via AAA/RADIUS**

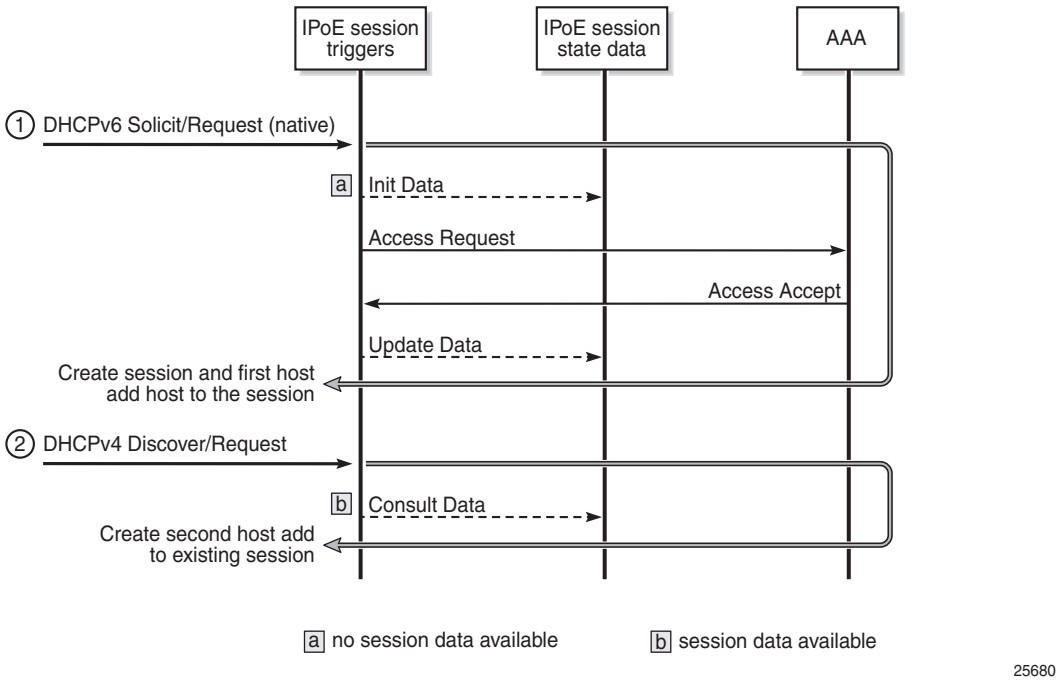

An Access-Request message is sent to the AAA/RADIUS server for authentication. On successful authentication, the AAA/RADIUS returns the ESM data (subscriber ID, sub-profile, SLA profile, and so on) in an Access-Accept message.

The Access-Request message optionally includes the host ID or ASID.

The behavior when authenticating through LUDB is similar.

IPoE Session Re-Authentication

IPoE sessions can be re-authenticated, meaning that the ESM data is fetched again from the ESM data source (LUDB, RADIUS, and so on), and controlled through the min-auth-interval timer.

If a host renewal packet is received while the min-auth-interval timer is running for the corresponding IPoE session, the cached ESM data is used (with one exception: see forced authentication). If a host renewal packet is received when this timer has expired, re-authentication is performed. If re-authentication fails, the host renewal packet is dropped; see Figure 156.

IPoE session re-authentication is configured in the ipoe-session context using following command:

```
min-auth-interval ?
  - min-auth-interval [days <days>] [hrs <hours>] [min <minutes>] [sec <seconds>]
  - min-auth-interval infinite
  - no min-auth-interval
 <days>               : [0..365]
 <hours>              : [0..23]
 <minutes>            : [0..59]
 <seconds>            : [0..59]
 <infinite>           : keyword
```

By default, re-authentication is disabled by having the min-auth-interval set to infinite. Setting the min-auth-interval to zero will lead to every single message (DORA, SARR) triggering re-authentication, but that is not recommended.

IPoE session re-authentication can be used to implement dynamic policy changes. For alternatives also implementing dynamic changes, see the Mid-Session Changes section.

For IPoE sessions, the re-authentication option in the RADIUS authentication-policy context is ignored.

IPoE Session Forced Authentication

Forced authentication means that the ESM data is fetched again from the ESM data source, regardless of the value of the re-authentication timer.

By default, forced authentication occurs when the CID or RID in the trigger packet has changed value, but this behavior can be disabled.

An absent or empty CID or RID is not considered as a change.

Forced authentication is configured in the ipoe-session context using following command:

```
force-auth ?
  - force-auth [cid-change] [rid-change]
  - force-auth disabled
  - no force-auth
 <cid-change>         : keyword - ignore min-auth-interval when cid changed
 <rid_change>         : keyword - ignore min-auth-interval when rid changed
 <disabled>           : keyword - never ignore min-auth-interval
```

IPoE Session Deletion

When the last host associated with an IPoE session is deleted, the IPoE session is deleted.

IPoE sessions are forcibly deleted in following situations:

- group-interface ipoe-session shutdown
- clear service <id> ipoe-session
- session timeout
- RADIUS disconnect or Diameter Gx RAR
- Credit Control (out-of-credit)

In all these cases, all hosts belonging to that session are deleted, with one exception. When the SLAAC inactivity-timer expires, only the corresponding SLAAC host is deleted, not the remaining hosts. When this SLAAC host is the last host of the IPoE session, the IPoE session is deleted.

The IPoE session-timeout is configured in the ipoe-session-policy:

```
configure
    subscriber-mgmt
        ipoe-session-policy <pol-name>
            session-timeout <timeout>
```

The timeout value ranges from 1 to 31104000 seconds (360 days). By default, no session-timeout is specified.

When RADIUS or Diameter Gx returns the Session-Timeout [27] or the Alc-Relative-Session-Timeout [26-6527-160] attributes, these values are used and the behavior is the same as for PPP sessions.

When no Session-Timeout or Alc-Relative-Session-Timeout attribute is returned by RADIUS, the session-timeout as configured in the ipoe-session-policy is used.

A RADIUS disconnect message, even when targeted at a single host, will also lead to the deletion of the entire IPoE session including all associated hosts.

A shutdown in the ipoe-session context of the group interface results in the deletion of all its IPoE sessions and associated hosts.

Session and Host limits

The number of IPoE sessions on a group interface and on a SAP can be limited:

```
configure service ies|vprn <service-id>
  subscriber-interface <ip-int-name>
     group-interface <ip-int-name>
       ipoe-session
         session-limit     [1..max*]
         sap-session-limit [1..max*]
```

The default values for the session-limit and the sap-session-limit are unlimited (no session-limit) and 1, respectively.

For retail services, the IPoE session-limit is configured at the linked subscriber interface level:

```
configure service vprn <retail-service-id>
    subscriber-interface <RT-ip-int-name> fwd-service <WS-service-id>
                              fwd-subscriber-interface <WS-ip-int-name>
        session-limit     [1..max*]
```

The default session-limit is unlimited.

Additionally, host limits can be imposed through the SLA profile:

```
configure subscriber-mgmt
   sla-profile <subscriber-profile-name>
     host-limits
       ipv4-arp            - Maximum number of IPv4 ARP hosts
       ipv4-dhcp           - Maximum number of IPv4 DHCP hosts
       ipv4-overall        - Maximum number of IPv4 hosts
       ipv4-ppp            - Maximum number of IPv4 PPP hosts
       ipv6-overall        - Maximum number of IPv6 hosts
       ipv6-pd-ipoe-dhcp   - Maximum number of IPv6-PD IPOE DHCP hosts
       ipv6-pd-overall     - Maximum number of IPv6-PD hosts
       ipv6-pd-ppp-dhcp    - Maximum number of IPv6-PD PPP DHCP hosts
       ipv6-wan-ipoe-dhcp  - Maximum number of IPv6-Wan IPOE DHCP hosts
       ipv6-wan-ipoe-slaac - Maximum number of IPv6-Wan IPOE SLAAC hosts
       ipv6-wan-overall    - Maximum number of IPv6-Wan hosts
       ipv6-wan-ppp-dhcp   - Maximum number of IPv6-Wan PPP DHCP hosts
       ipv6-wan-ppp-slaac  - Maximum number of IPv6-Wan PPP SLAAC hosts
       lac-overall         - Maximum number of L2TP LAC hosts
       overall             - Maximum number of hosts
       remove-oldest       - Remove oldest
     exit
```

See the Wholesale/Retail section for more information about limits in a wholesale/ retail configuration.

RADIUS Session Accounting

As well as queue instance and host accounting, RADIUS session accounting can be enabled for IPoE sessions.

Usually, a RADIUS Accounting-Start message is sent when the first host is associated with an IPoE session. Regular and triggered accounting Interim-Update messages are sent during the IPoE session. An Accounting-Stop message is sent when the last host is deleted from the session.

Session accounting is configured in the RADIUS accounting policy, and can be set to the following values:

- session-accounting
- session-accounting interim-update
- session-accounting host-update
- session-accounting interim-update host-update

Plain session accounting sends start and stop messages. The RADIUS accounting server is informed about the start and the stop time of the session, but no counters are maintained. This implements time-based accounting.

The interim-update option additionally sends interim-update messages, so that the RADIUS accounting server maintains counters. This implements volume accounting.

The host-update option additionally sends host up/down event messages, so that the RADIUS accounting server keeps track of host creation and deletion events.

The combination of the interim-update and host-update options allows the RADIUS accounting server to track all changes.

Mid-Session Changes

Mid-session changes, such as those initiated via RADIUS CoA or Diameter Gx RAR messages, are applied to all hosts associated with the IPoE session. There is no way to update a single host of an 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 hosts of the session are targeted and the session state data is updated.

Mid-session changes also can be applied manually, using the following command:

```
# 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>]
```

The tools commands eval-lease-state and eval-slaac-host are blocked when the host is part of a session.

IPoE session re-authentication can also lead to dynamic policy changes.

Subscriber Host Connectivity Verification

Subscriber host connectivity verification (SHCV) can be enabled for hosts associated with an IPoE session.

When a single host fails and stops responding to the SHCV messages, that host is deleted without affecting the other hosts that are part of the session. When the last host fails, the session is deleted.

IA-PD managed routes are not subject to SHCV, and cannot be removed because of SHCV.

Dual Homing

IPoE sessions are supported in a dual-homed environment, where Multichassis Synchronization (MCS) and Subscriber Routed Redundancy Protocol (SRRP) are active.

MCS ensures that the IPoE session data is synchronized between the BNG pair.

Wholesale/Retail

IPoE sessions are supported in single-homed and dual-homed wholesale/retail environments.

The wholesale IPoE session limit is configured at the group interface level, and the retail IPoE session limit is configured at the linked subscriber interface level:

```
configure service vprn <retail-service-id>
    subscriber-interface <RT-ip-int-name> fwd-service <WS-service-id>
                                          fwd-subscriber-interface <WS-ip-int-name>
        ipoe-session
            session-limit     [1..max*]  #default unlimited
```

In IPoE, private-retail-subnets only apply to IPoEv6 in single-homed scenarios (no MCS). Therefore, the use cases for private-retail-subnets in combination with IPoE sessions are limited.

Practical Considerations

The rules for configuring authentication for regular, capture, and managed SAPs, also apply to IPoE sessions; see Figure 158:

- If an authentication policy is applied at capture-SAP or group-interface level, this policy has priority, regardless of whether, or in which other sub-contexts, an LUDB is assigned. Therefore, for an LUDB to provide ESM data, no authentication policy may be applied at capture-SAP or group interface level.
- If an LUDB is applied in the ipoe-session context of a group interface or capture SAP, the LUDBs assigned in the dhcp, dhcp6, and router-solicit related contexts of the same group interface or capture SAP are ignored.

Figure 158 **Configuring IPoE session authentication**

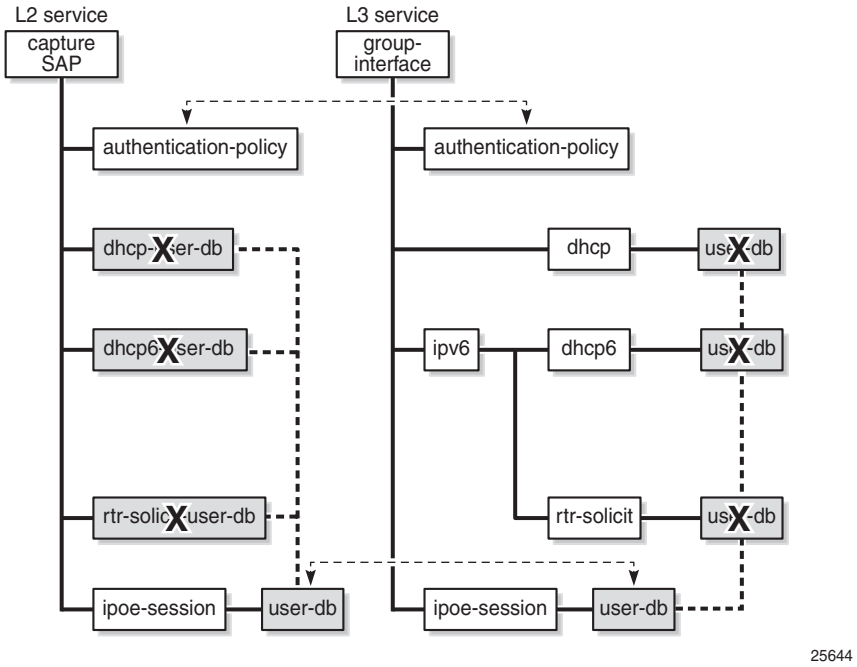

When the AAA/RADIUS server referenced from the authentication policy is not available, SR OS can rely on a fallback LUDB, if configured; see the LUDB for ESM chapter for more information.

Be aware of the following:

- Static hosts can be configured on a group interface with IPoE sessions enabled. A static host will not be associated with an IPoE session.
- ARP hosts are not supported in an IPoE session and cannot be instantiated on a group interface with IPoE sessions enabled.
- Up to sixteen framed-routes and sixteen framed-IPv6-routes can be associated with an IPoE session.
- Python-based subscriber identification based on the DHCPv4 Ack message is ignored when IPoE sessions are enabled.

Configuration

Baseline configuration

Figure 6 shows the baseline configuration used for the examples in this chapter, excluding the ipoe-session configurations. These will be added later in this chapter.

The first example uses LUDB authentication, and the second example uses AAA/RADIUS authentication. As alternatives, AAA/NASREQ authentication or AAA/Gx authentication can be used.

Figure 159 **Baseline configuration**

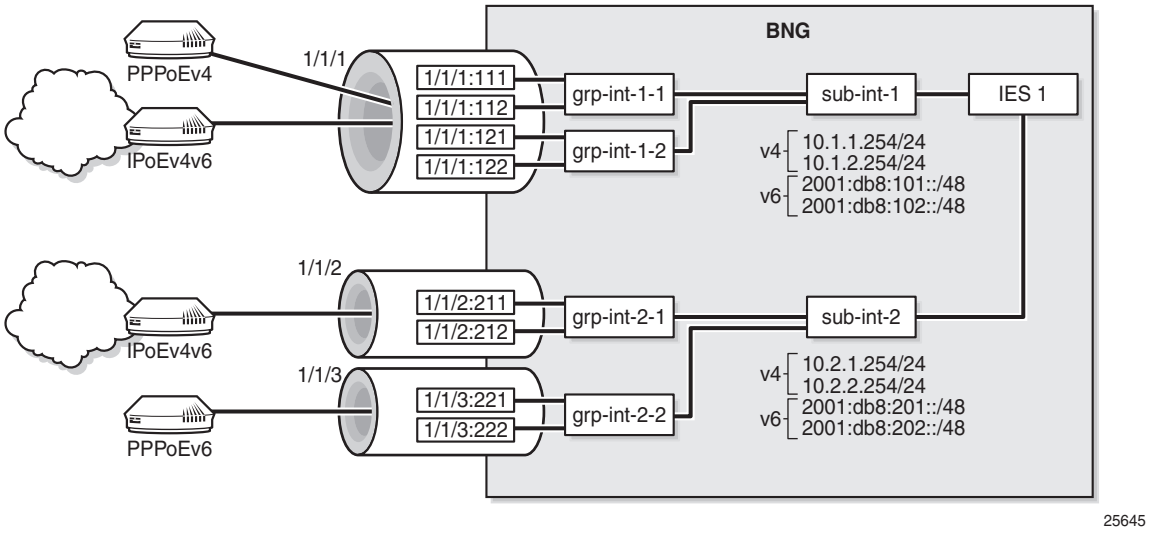

The following partial configuration applies to IES-1. This service is provisioned with ESM on all of its SAPs, and supports proxy and relay scenarios on all group interfaces for both IPv4 and IPv6. Only the part relevant to subscriber interface *sub-int-1* and group interface *grp-int-1-1* is shown. The configurations for the other subscriber and group interfaces are similar. See the ESM Basics and Routed CO chapters for more information.

```
configure
    service
        ies 1 customer 1 create
            subscriber-interface "sub-int-1" create
                address 10.1.1.254/24
                --- snipped ---
                ipv6
                    delegated-prefix-len 56
                    link-local-address fe80::ea:4b:ff
                    subscriber-prefixes
                        prefix 2001:db8:101::/48 wan-host
                        prefix 2001:db8:f101::/48 pd
                        --- snipped ---
                    exit
                exit
                group-interface grp-int-1-1
                    ipv6
                        router-advertisements
                            prefix-options
                                autonomous
                            exit
                            no shutdown
                        exit
                        dhcp6
                            proxy-server
                                client-applications dhcp ppp
```

```
                no shutdown
            exit
            relay
                link-address 2001:db8:101::1
                server 2001:db8::11
                client-applications dhcp ppp
                no shutdown
            exit
        exit
        router-solicit
            no shutdown
        exit
    exit
    local-address-assignment
        ipv6
            client-application ipoe-slaac
            server "dhcp6-srv"
        exit
        no shutdown
    exit
    arp-populate
    dhcp
        proxy-server
            emulated-server 10.1.1.254
            no shutdown
        exit
        server 10.11.11.1
        trusted
        lease-populate 1000
        client-applications dhcp ppp
        gi-address 10.1.1.254
        no shutdown
    exit
    --- snipped ---
    sap 1/1/1:111 create
        sub-sla-mgmt
            def-sub-profile "sub-prof-1"
            def-sla-profile "sla-prof-1"
            sub-ident-policy "sub-id-pol-1"
            multi-sub-sap
            no shutdown
        exit
    exit
    --- snipped ---
```

No DHCPv4 or DHCPv6 relay options are defined.

Troubleshooting

The syntax to show the active IPoE sessions is as follows:

```
show service id <service-id> ipoe session ?
  - 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>] [wholesaler <service-id>]
  - 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>]
```

The following show commands have been extended, so that session filtering is available:

```
show service id <svc-id> dhcp lease-state ?
  - 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]

show service id <svc-id> dhcp6 lease-state ?
  - lease-state [detail] [wholesaler <service-id>] [session {none|ipoe|ppp}]
  - lease-state [detail] interface <interface-name> [wholesaler <service-
id>] [session {none|ipoe|ppp}]
  - lease-state [detail] ipv6-address <ipv6-prefix[/prefix-
length]> [wholesaler <service-id>] [session {none|ipoe|ppp}]
  - lease-state [detail] mac <ieee-address> [wholesaler <service-
id>] [session {none|ipoe|ppp}]

show service id <svc-id> slaac host ?
  - host [detail] [wholesaler <service-id>] [session {none|ipoe|ppp}]
  - host interface <interface-name> [detail] [wholesaler <service-
id>] [session {none|ipoe|ppp}]
  - host mac <ieee-address> [detail] [wholesaler <service-
id>] [session {none|ipoe|ppp}]
  - host ipv6-address <ipv6-prefix> [detail] [wholesaler <service-
id>] [session {none|ipoe|ppp}]
  - host sap <sap-id> [detail] [wholesaler <service-id>] [session {none|ipoe|ppp}]
```

The following debug configuration is used for demonstration and troubleshooting purposes:

```
debug
    router "Base"
        ip
            dhcp
                detail-level medium
                mode egr-ingr-and-dropped
            exit
            dhcp6
                mode egr-ingr-and-dropped
                detail-level medium
            exit
        exit
        radius
            packet-type authentication accounting coa
            detail-level medium
        exit
    exit
```

```
    subscriber-mgmt
        local-user-db "ludb-1"
            detail all
        exit
    exit
exit
```

IPoE session failure events are also issued to log-id 99:

```
*A:BNG-1# show log event-control "svcmgr"
===============================================================================
Log Events
===============================================================================
Application
 ID#    Event Name                          P   g/s     Logged     Dropped
-------------------------------------------------------------------------------
   2011 svcTlsMacPinningViolation           WA  thr          0           0
   --- snipped ---
   2554 tmnxSubIpoeInvalidSessionKey        WA  thr         30           0
   2555 tmnxSubIpoeInvalidCidRidChange      WA  thr          0           0
   2556 tmnxSubIpoeSessionLimitReached      WA  thr          0           0
   2557 tmnxSubIpoePersistenceRecovery      WA  thr          0           0
   2559 tmnxSubIpoeMigrHostDeleted          WA  thr          0           0
===============================================================================
*A:BNG-1#
```

IPoE Session Authentication through LUDB

The LUDB *ludb-1* uses the MAC address for host matching, and is defined as follows:

```
configure
    subscriber-mgmt
        local-user-db "ludb-1" create
            description "example user-db"
            ipoe
                match-list mac
                host "entry-01" create
                    host-identification
                        mac 00:00:00:00:00:01
                    exit
                    address pool "pool4-1"
                    identification-strings 254 create
                        subscriber-id "sub-11"
                        sla-profile-string "sla-profile-1"
                        sub-profile-string "sub-profile-1"
                    exit
                    ipv6-slaac-prefix-pool "pool6-1"
                    ipv6-wan-address-pool "pool6-1"
                    no shutdown
                exit
                --- snipped ---
```

This LUDB is then applied to the group interface in the ipoe-session context:

```
configure
    service
        ies 1 customer 1 create
            subscriber-interface "sub-int-1"
                group-interface "grp-int-1-1"
                    ipoe-session
                        description "ipoe-sessions with LUDB"
                        ipoe-session-policy "sespol-sap-mac"
                        sap-session-limit 100
                        session-limit 500
                        user-db "ludb-1"
                        no shutdown
                    exit
```

The LUDB applied in the ipoe-session context takes priority over LUDBs applied in the *dhcp*, *ipv6 dhcpv6*, and *ipv6* router-solicit contexts for a Layer 3 service.

The IPoE session policy sespol-sap-mac used in this example is defined as follows:

```
configure
    subscriber-mgmt
        ipoe-session-policy "sespol-sap-mac" create
            description "plain ipoe session policy, sap-mac key"
            session-key sap mac
            no session-timeout
        exit
```

Debug

The following debug trace appears when the user with MAC address 00:00:00:00:00:01 first connects using DHCPv4 and subsequently connects using SLAAC and DHCPv6 (IANA), without disconnecting DHCPv4.

Messages 1 through 9 show the message sequence for DHCPv4 (DORA). Messages 11 through 22 show the message sequence for DHCPv6 (SARR). Message 10 and 23 are the router solicitation and advertisement messages. Therefore, three hosts are created.

The LUDB is accessed just once; immediately after the DHCPv4 Discover message:

```
1 2016/02/03 13:51:28.15 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 6 (grp-int-1-1),
   received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:111) Port 67
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:00:00:00:00:01    xid: 0x1
   DHCP options:
```

```
   [53] Message type: Discover
   [255] End
"
2 2016/02/03 13:51:28.15 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  mac: 00:00:00:00:00:01
  Host entry-01 found in user data base ludb-1"
3 2016/02/03 13:51:28.15 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
   transmitted DHCP Boot Request to 10.11.11.1 Port 67
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 10.1.1.254
   chaddr: 00:00:00:00:00:01    xid: 0x1
   DHCP options:
   [53] Message type: Discover
   [255] End
"
--- snipped ---
9 2016/02/03 13:51:28.16 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 6 (grp-int-1-1),
   transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:111) Port 68
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.1.1.60
   siaddr: 10.11.11.1        giaddr: 10.1.1.254
   chaddr: 00:00:00:00:00:01    xid: 0x1
   DHCP options:
   [53] Message type: Ack
   [54] DHCP server addr: 10.11.11.1
   [51] Lease time: 900
   [1] Subnet mask: 255.255.255.0
   [3] Router: 10.1.1.254
   [255] End
"
10 2016/02/03 13:51:40.77 CET MINOR: DEBUG #2001 Base TIP
"TIP: ICMP6_PKT
ICMP6 ingressing on grp-int-1-1 (Base):
   fe80::200:ff:fe00:1 -> ff02::2
   Type: Router Solicitation (133)
   Code: No Code (0)
      Option  : Src Link Layer Addr 00:00:00:00:00:01
"
11 2016/02/03 13:51:40.78 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Incoming DHCP6 Msg : SOLICIT (1)
   on itf grp-int-1-1
     Trans Id : 0x411fc8
     Option : CLIENTID (1), Length : 14
       LLT : HwTyp=0001,T=507311564,LL=000000000001
       000100011e3cf5cc000000000001
     Option : IA_NA (3), Length : 12
       IAID : 2
       Time1: 0 seconds
       Time2: 0 seconds
     Option : ORO (6), Length : 2
       Requested Option : DNS_NAME_SRVR (23)
"
```

```
--- snipped ---
22 2016/02/03 13:51:40.92 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Outgoing DHCP6 Msg : REPLY (7)
   to itf grp-int-1-1
     Trans Id : 0xe15ddf
     Option : SERVERID (2), Length : 10
       LL  : HwTyp=0001,LL=ea4bff000000
       00030001ea4bff000000
     Option : CLIENTID (1), Length : 14
       LLT : HwTyp=0001,T=507311564,LL=000000000001
       000100011e3cf5cc000000000001
     Option : IA_NA (3), Length : 40
       IAID : 2
       Time1: 1800 seconds
       Time2: 2880 seconds
       Option : IAADDR (5), Length : 24
         Address : 2001:db8:101:c::1
         Preferred Lifetime : 3600 seconds
         Valid Lifetime     : 86400 seconds
     Option : DNS_NAME_SRVR (23), Length : 16
       Server : 2001:db8::1:1:1:1
"
23 2016/02/03 13:51:41.91 CET MINOR: DEBUG #2001 Base TIP
"TIP: ICMP6_PKT
ICMP6 egressing on grp-int-1-1 (Base):
   fe80::ea:4b:ff -> fe80::200:ff:fe00:1
   Type: Router Advertisement (134)
   Code: No Code (0)
      Hop Limit     : 64
      Flags         :
      Retrans Time  : 0
      Def Life Time : 4500
      Reachable Time: 0
      Option  : Src Link Layer Addr 00:00:5e:00:01:01
      Option  : Prefix         : 2001:db8:101:d::/64
                Flags          : On Link Autoconfig
                Valid Life Time: 86400
                Pref  Life Time: 3600
"
```

Verification

The following shows the IPoE session for MAC address 00:00:00:00:00:01:

```
*A:BNG-1# show service id 1 ipoe session mac 00:00:00:00:00:01
===============================================================================
IPoE sessions for svc-id 1
===============================================================================
Sap Id                          Mac Address         Up Time         MC-Stdby
    Subscriber-Id
        [CircuitID] | [RemoteID]
-------------------------------------------------------------------------------
1/1/1:111                       00:00:00:00:00:01   0d 00:10:18
    sub-11
```

```
-------------------------------------------------------------------------------
CID | RID displayed when included in session-key
Number of sessions : 1
===============================================================================
*A:BNG-1#
```

The IPoE session details for MAC address 00:00:00:00:00:01 are shown using the following command. The session time left is undefined (N/A), because no IPoE session-timeout is defined in the IPoE session policy. Re-authentication does not apply, so the minimum authentication interval is infinite (N/A).

```
*A:BNG-1# show service id 1 ipoe session mac 00:00:00:00:00:01 detail
===============================================================================
IPoE sessions for service 1
===============================================================================
SAP                     : 1/1/1:111
Mac Address             : 00:00:00:00:00:01
Circuit-Id              :
Remote-Id               :
Session Key             : sap-mac
MC-Standby              : No
Subscriber-interface    : sub-int-1
Group-interface         : grp-int-1-1
Termination Type        : local
Up Time                 : 0d 00:09:42
Session Time Left       : N/A
Last Auth Time          : 02/03/2016 13:51:29
Min Auth Intvl (left)   : infinite (N/A)
Persistence Key         : N/A
Subscriber              : "sub-11"
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         : "EA4BFF000001A656B1F7D0"
Sap-Session-Index       : 1
IP Address              : 10.1.1.60/24
IP Origin               : DHCP
Primary DNS             : N/A
Secondary DNS           : N/A
Primary NBNS            : N/A
Secondary NBNS          : N/A
Address-Pool            : pool4-1
IPv6 Prefix             : 2001:db8:101:d::/64
IPv6 Prefix Origin      : LclPool
IPv6 Prefix Pool        : "pool6-1"
IPv6 Del.Pfx.           : N/A
IPv6 Del.Pfx. Origin    : None
IPv6 Del.Pfx. Pool      : ""
IPv6 Address            : 2001:db8:101:c::1
IPv6 Address Origin     : DHCP
IPv6 Address Pool       : "pool6-1"
Primary IPv6 DNS        : 2001:db8::1:1:1:1
Secondary IPv6 DNS      : N/A
Radius Session-TO       : N/A
```

```
Radius Class            :
Radius User-Name        :
-------------------------------------------------------------------------------
Number of sessions : 1
===============================================================================
*A:BNG-1#
```

The following command shows the subscriber hosts for this MAC address:

```
*A:BNG-1# show service id 1 subscriber-hosts mac 00:00:00:00:00:01
============================================================
Subscriber Host table
============================================================
Sap                    Subscriber
  IP Address
    MAC Address           PPPoE-SID Origin       Fwding State
------------------------------------------------------------
1/1/1:111              sub-11
  10.1.1.60
    00:00:00:00:00:01    N/A       DHCP         Fwding
1/1/1:111              sub-11
  2001:db8:101:c::1/128
    00:00:00:00:00:01    N/A       IPoE-DHCP6   Fwding
1/1/1:111              sub-11
  2001:db8:101:d::/64
    00:00:00:00:00:01    N/A       IPoE-SLAAC   Fwding
------------------------------------------------------------
Number of subscriber hosts : 3
============================================================
*A:BNG-1#
```

The following commands show the corresponding dhcp and dhcp6 lease-states:

```
*A:BNG-1# show service id 1 dhcp lease-state session ipoe
===============================================================================
DHCP lease state table, service 1
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LeaseTime  Origin   Stdby
-------------------------------------------------------------------------------
10.1.1.60       00:00:00:00:00:01 1/1/1:111           00h13m17s  DHCP
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
*A:BNG-1#
*A:BNG-1# show service id 1 dhcp6 lease-state session ipoe
===============================================================================
DHCP lease state table, service 1
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LeaseTime  Origin   Stdby
-------------------------------------------------------------------------------
2001:db8:101:c::1/128
                00:00:00:00:00:01 1/1/1:111           23h59m02s  DHCP
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
```

```
*A:BNG-1#
```

IPoE Session Authentication through AAA/RADIUS

The FreeRADIUS server users file contains following data for the connecting device:

```
00:00:00:00:00:01    Cleartext-Password := "spasswd"
                     Alc-Subsc-ID-Str = "ipoe-%{User-name}",
                     Alc-Subsc-Prof-Str = "sub-prof-1",
                     Alc-SLA-Prof-Str = "sla-prof-1",
                     Framed-Pool = "pool4-1",
                     Framed-Ipv6-Pool = "pool6-1",
                     Alc-Delegated-IPv6-Pool = "pool6-1",
                     Alc-Relative-Session-Timeout = "300"
```

The authentication policy *radius-pol* used in this example is defined as follows:

```
configure
    subscriber-mgmt
        authentication-policy "radius-pol"
            password letmein
            include-radius-attribute
                acct-session-id
                circuit-id
                sap-session-index
            exit
            radius-server-policy "rsp-1"
        exit
```

The IPoE session policy used in this example is defined as follows. The key now also includes the circuit ID.

```
configure
    subscriber-mgmt
        ipoe-session-policy "sespol-sap-mac-cid" create
            description "key also including cid now"
            session-key sap mac cid
            no session-timeout
        exit
```

The authentication and IPoE session policies are then applied to the group interface g*rp-int-1-1*:

```
configure
   service
        ies 1 customer 1 create
            subscriber-interface "sub-int-1"
                group-interface "grp-int-1-1"
                    authentication-policy "radius-pol"
                    ipoe-session
                        ipoe-session-policy "sespol-sap-mac-cid"
                        sap-session-limit 100
```

```
            session-limit 100
            no shutdown
        exit
```

The authentication policy takes precedence over any LUDB applied in one of the subcontexts of that group interface.

Debug

The following debug trace appears when the user with MAC address 00:00:00:00:00:01 first connects using DHCPv6 (IA_NA and IA_PD) and subsequently connects using DHCPv4.

The ESM data source is accessed just once, immediately after the DHCPv6 Solicit message. Because LDRA is active, the Solicit message is embedded in the Relay Forward message. The RADIUS server is sent an Access-Request message including the circuit ID, and returns an Access-Accept message including the Alc-Relative-Session-Timeout attribute (message 3). Message 14 is the final IPv6 Reply message containing both the IA_NA address and the IA_PD prefix used by the CPE. Messages 15 through 22 are the DHCPv4 DORA messages.

The initial DHCPv6 Solicit message and the initial DHCPv4 Discover message contain the same interface ID/CID (11), which is why no re-authentication is triggered. The IPoE session is deleted when the RADIUS-provided session timer expires, so the BNG releases both the IPv4 and the IPv6 address (messages 23 through 31).

```
1 2016/02/03 14:51:27.54 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Incoming DHCP6 Msg : RELAY_FORW (12)
   on itf grp-int-1-1
     Hop Count : 0
     Link Addr : ::
     Peer Addr : fe80::200:ff:fe00:1
     Option : RELAY_MSG (9), Length : 60
       Msg Type : SOLICIT (1)
       Trans Id : 0x89ade3
       Option : CLIENTID (1), Length : 14
         LLT : HwTyp=0001,T=507311564,LL=000000000001
         000100011e3cf5cc000000000001
       Option : IA_PD (25), Length : 12
         IAID : 1
         Time1: 0 seconds
         Time2: 0 seconds
       Option : IA_NA (3), Length : 12
         IAID : 2
         Time1: 0 seconds
         Time2: 0 seconds
       Option : ORO (6), Length : 2
         Requested Option : DNS_NAME_SRVR (23)
```

```
     Option : INTERFACE_ID (18), Length : 2
       Interface Id : 3131 (11)
"
2 2016/02/03 14:51:27.54 CET MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.16.1.1:1812 id 165 len 109 vrid 1 pol rsp-1
    USER NAME [1] 17 00:00:00:00:00:01
    PASSWORD [2] 16 8DF.2ZKk.XRvmblEXcKEOk
    NAS IP ADDRESS [4] 4 192.0.2.1
    VSA [26] 4 DSL(3561)
      AGENT CIRCUIT ID [1] 2 11
    SESSION ID [44] 22 EA4BFF000001C356B205DF
    VSA [26] 6 Alcatel(6527)
      SAP SESSION INDEX [180] 4 1
"
3 2016/02/03 14:51:27.54 CET MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 165 len 131 from 172.16.1.1:1812 vrid 1 pol rsp-1
    VSA [26] 24 Alcatel(6527)
      SUBSC ID STR [11] 22 ipoe-00:00:00:00:00:01
    VSA [26] 12 Alcatel(6527)
      SUBSC PROF STR [12] 10 sub-prof-1
    VSA [26] 12 Alcatel(6527)
      SLA PROF STR [13] 10 sla-prof-1
    FRAMED POOL [88] 7 pool4-1
    FRAMED IPV6 POOL [100] 7 pool6-1
    VSA [26] 9 Alcatel(6527)
      DELEGATED IPV6 POOL [131] 7 pool6-1
    VSA [26] 6 Alcatel(6527)
      RELATIVE SESSION TIMEOUT [160] 4 300
"
4 2016/02/03 14:51:27.55 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Outgoing DHCP6 Msg : RELAY_FORW (12)
   to itf int-DHCP
     Hop Count : 1
     Link Addr : 2001:db8:101::1
     Peer Addr : fe80::200:ff:fe00:1
     Option : RELAY_MSG (9), Length : 104
       Msg Type : RELAY_FORW (12)
       Hop Count : 0
       Link Addr : ::
       Peer Addr : fe80::200:ff:fe00:1
       Option : RELAY_MSG (9), Length : 60
         Msg Type : SOLICIT (1)
         Trans Id : 0x89ade3
         --- snipped ---
     Option : VENDOR_OPTS (17), Length : 36
       Enterprise : 0000197f
       Option : WAN_POOL (1), Length : 7
         pool6-1
       Option : PFX_POOL (2), Length : 7
         pool6-1
       Option : PFX_LEN (3), Length : 1
         56
       Option : RESERVED_NA_LEN (4), Length : 1
         64
"
--- snipped ---
```

```
14 2016/02/03 14:51:27.67 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Outgoing DHCP6 Msg : RELAY_REPLY (13)
   to itf grp-int-1-1
     Hop Count : 0
     Link Addr : ::
     Peer Addr : fe80::200:ff:fe00:1
     Option : RELAY_MSG (9), Length : 145
       Msg Type : REPLY (7)
       Trans Id : 0x4c546f
       --- snipped ---
       Option : IA_PD (25), Length : 41
         IAID : 1
         Time1: 1800 seconds
         Time2: 2880 seconds
         Option : IAPREFIX (26), Length : 25
           Prefix : 2001:db8:f101:700::/56
           Preferred Lifetime : 3600 seconds
           Valid Lifetime     : 86400 seconds
       Option : IA_NA (3), Length : 40
         IAID : 2
         Time1: 1800 seconds
         Time2: 2880 seconds
         Option : IAADDR (5), Length : 24
           Address : 2001:db8:101:c::1
           Preferred Lifetime : 3600 seconds
           Valid Lifetime     : 86400 seconds
       Option : DNS_NAME_SRVR (23), Length : 16
         Server : 2001:db8::1:1:1:1
     Option : INTERFACE_ID (18), Length : 2
       Interface Id : 3131 (11)
"
15 2016/02/03 14:51:34.59 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 6 (grp-int-1-1),
   received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:111) Port 67
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:00:00:00:00:01    xid: 0x1
   DHCP options:
   [82] Relay agent information: len = 4
      [1] Circuit-id: 11
   [53] Message type: Discover
   [255] End
"
--- snipped ---
22 2016/02/03 14:51:34.73 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 6 (grp-int-1-1),
   transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:111) Port 68
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.1.1.64
   siaddr: 10.11.11.1        giaddr: 10.1.1.254
   chaddr: 00:00:00:00:00:01    xid: 0x1
   DHCP options:
   [82] Relay agent information: len = 4
      [1] Circuit-id: 11
   [53] Message type: Ack
```

```
   [54] DHCP server addr: 10.11.11.1
   [51] Lease time: 900
   [1] Subnet mask: 255.255.255.0
   [3] Router: 10.1.1.254
   [255] End
"
23 2016/02/03 14:56:26.55 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Outgoing DHCP6 Msg : RELAY_FORW (12)
   to itf int-DHCP
     Hop Count : 0
     Link Addr : 2001:db8:101::1
     Peer Addr : fe80::200:ff:fe00:1
     Option : RELAY_MSG (9), Length : 80
       Msg Type : RELEASE (8)
       Trans Id : 0x000000
       --- snipped ---
     Option : INTERFACE_ID (18), Length : 22
       Interface Id : 5f746d6e785f696e7465726e616c5f636c65616e7570 ( snipped )
"
--- snipped ---
26 2016/02/03 14:56:26.55 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
   transmitted DHCP Boot Request to 10.11.11.1 Port 68
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 10.1.1.64         yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:00:00:00:00:01    xid: 0x0
   DHCP options:
   [53] Message type: Release
   [54] DHCP server addr: 10.11.11.1
   [255] End
"
--- snipped ---
31 2016/02/03 14:56:26.56 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Incoming DHCP6 Msg : RELAY_REPLY (13)
   on itf int-DHCP
     Hop Count : 0
     Link Addr : 2001:db8:101::1
     Peer Addr : fe80::200:ff:fe00:1
     Option : RELAY_MSG (9), Length : 89
       Msg Type : REPLY (7)
       --- snipped ---
       Option : IA_PD (25), Length : 49
         IAID : 1
         Time1: 0 seconds
         Time2: 0 seconds
         Option : STATUS_CODE (13), Length : 33
           Status : SUCCESS (0)
                    All prefixes have been released
     Option : INTERFACE_ID (18), Length : 22
       Interface Id : 5f746d6e785f696e7465726e616c5f636c65616e7570 ( snipped )
"
```

Verification

The following command shows the session details for MAC address 00:00:00:00:00:01. The key now includes the circuit ID (11), and the session timer is running. The RADIUS-provided session timeout is 5 minutes.

```
*A:BNG-1# show service id 1 ipoe session mac 00:00:00:00:00:01 detail
===============================================================================
IPoE sessions for service 1
===============================================================================
SAP                     : 1/1/1:111
Mac Address             : 00:00:00:00:00:01
Circuit-Id              : 11
Remote-Id               :
Session Key             : sap-mac-cid
MC-Standby              : No
Subscriber-interface    : sub-int-1
Group-interface         : grp-int-1-1
Termination Type        : local
Up Time                 : 0d 00:03:13
Session Time Left       : 0d 00:01:47
Last Auth Time          : 02/03/2016 14:51:28
Min Auth Intvl (left)   : infinite (N/A)
Persistence Key         : N/A
Subscriber              : "ipoe-00:00:00:00:00:01"
Sub-Profile-String      : "sub-prof-1"
SLA-Profile-String      : "sla-prof-1"
ANCP-String             : ""
Int-Dest-Id             : ""
App-Profile-String      : ""
Category-Map-Name       : ""
Acct-Session-Id         : "EA4BFF000001C256B205DF"
Sap-Session-Index       : 1
IP Address              : 10.1.1.64/24
IP Origin               : DHCP
Primary DNS             : N/A
Secondary DNS           : N/A
Primary NBNS            : N/A
Secondary NBNS          : N/A
Address-Pool            : pool4-1
IPv6 Prefix             : N/A
IPv6 Prefix Origin      : None
IPv6 Prefix Pool        : ""
IPv6 Del.Pfx.           : 2001:db8:f101:700::/56
IPv6 Del.Pfx. Origin    : DHCP
IPv6 Del.Pfx. Pool      : "pool6-1"
IPv6 Address            : 2001:db8:101:c::1
IPv6 Address Origin     : DHCP
IPv6 Address Pool       : "pool6-1"
Primary IPv6 DNS        : N/A
Secondary IPv6 DNS      : N/A
Radius Session-TO       : 0d 00:05:00
Radius Class            :
Radius User-Name        : 00:00:00:00:00:01
-------------------------------------------------------------------------------
Number of sessions : 1
===============================================================================
```

```
*A:BNG-1#
```

Three hosts are created, as the following command shows:

```
*A:BNG-1# show service id 1 subscriber-hosts mac 00:00:00:00:00:01

===============================================================
Subscriber Host table
===============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID Origin       Fwding State
---------------------------------------------------------------
1/1/1:111              ipoe-00:00:00:00:00:01
  10.1.1.64
    00:00:00:00:00:01    N/A       DHCP         Fwding
1/1/1:111              ipoe-00:00:00:00:00:01
  2001:db8:101:c::1/128
    00:00:00:00:00:01    N/A       IPoE-DHCP6   Fwding
1/1/1:111              ipoe-00:00:00:00:00:01
  2001:db8:f101:700::/56
    00:00:00:00:00:01    N/A       IPoE-DHCP6   Fwding
---------------------------------------------------------------
Number of subscriber hosts : 3
===============================================================
*A:BNG-1#
```

After 300 seconds, the session is deleted:

```
*A:BNG-1# show service id 1 ipoe session detail
No entries found.
*A:BNG-1#
```

Conclusion

IPoE sessions offer ISPs a simplified way to manage dual-stack IPoE devices. IPoE sessions have features similar to PPP sessions, in terms of authentication, mid-session changes, and accounting. IPoE sessions can be used on regular, capture, and managed SAPs, and are supported in single- and dual-homed scenarios, including wholesale and retail configurations.

IPv4 DHCP Hosts

This chapter provides information about IPv4 DHCP host configurations.

Topics in this chapter include:

Applicability

This chapter is related to the use of IPv4 only, and was originally written for and tested on release 7.0.R6. The CLI now corresponds to release 13.0.R7.

Configuration and troubleshooting commands are given for Bridged CO and Routed CO scenarios.

Summary

In the Triple Play Service Delivery Architecture (TPSDA), a subscriber is defined as a collection of hosts pertaining to a single access connection (such as a DSL line) and identified by a subscriber identifier. A subscriber host is an end user terminal within the subscriber home (for example, a PC, set-top box, home gateway) that is identified in the network with a unique (IP address; MAC address) tuple for IPoE or (PPPoE session ID; MAC address) tuple for PPPoE.

Following IPv4 host types are distinguished:

- Static hosts
  - ip-mac
  - ip-only
- Dynamic hosts
  - ARP-host
  - DHCP-host

– PPPoE-host

This chapter provides configuration and troubleshooting commands for DHCP-hosts.

Knowledge of the Triple Play Service Delivery Architecture (TPSDA) concepts is assumed throughout this document.

Overview

The network topology for a Bridged CO environment is displayed in Figure 160 and for a Routed CO environment in Figure 161.

Figure 160 **Bridged CO Network Topology**

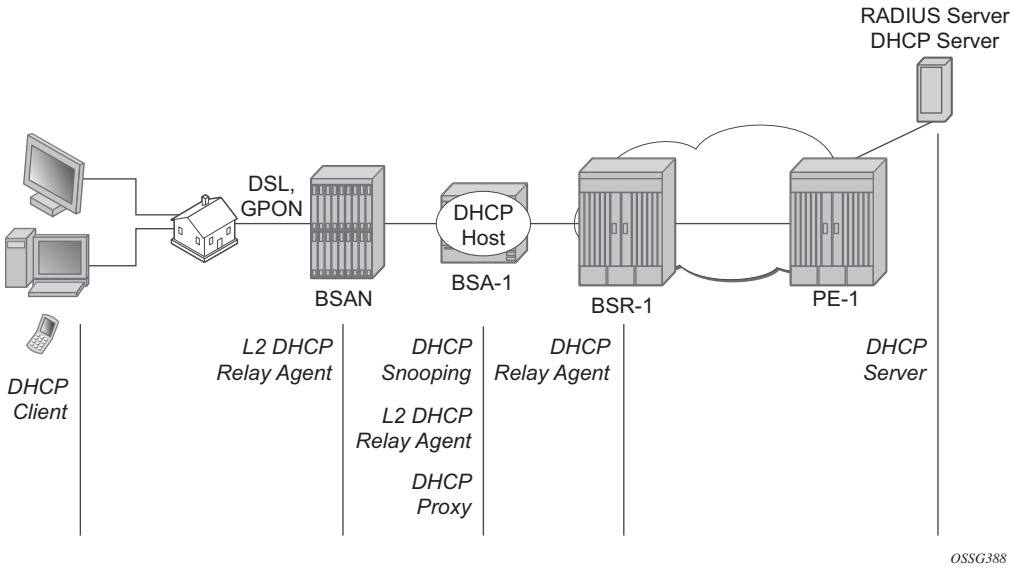

Figure 161 **Routed CO Network Topology**

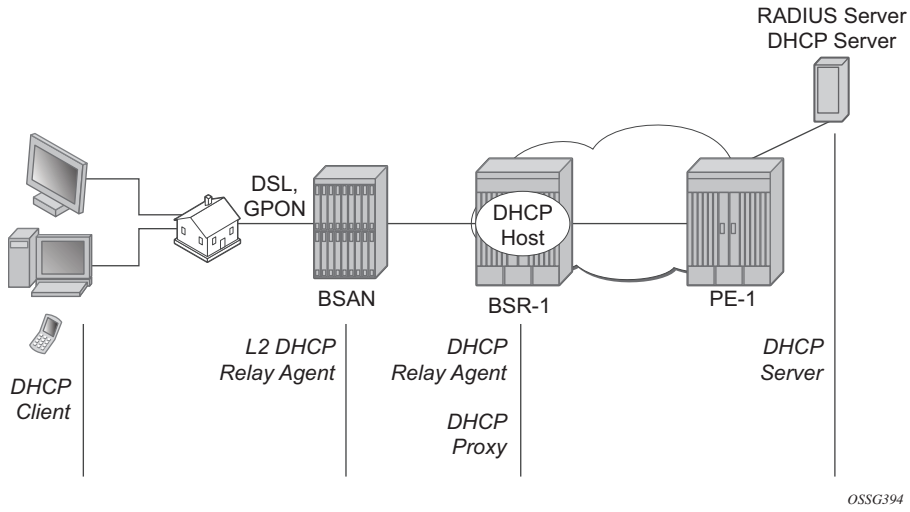

Following configuration tasks should be done first and are not detailed in this configuration note:

- Basic service router configurations such as system interface, IGP, MPLS, BGP.
- Bridged CO service topology: VPLS on BSA-1, terminated in a VPRN or IES service on BSR-1.
- Routed CO service topology: VPRN or IES service with subscriber and group interface on BSR-1.
- External DHCP server: server configuration and connectivity in the VPRN or base router instance.
- External RADIUS server: server configuration and connectivity in the VPRN or base router instance (Enhanced Subscriber Management (ESM) only).

This chapter focuses on DHCP hosts instantiated in a VPLS service on BSA-1 (Bridged CO) or in a VPRN service subscriber interface on BSR-1 (Routed CO). Note that in case of Routed CO, it is also possible to instantiate the DHCP hosts in the base routing instance using an IES service.

Most of the DHCP host functionality is available with Basic Subscriber Management (BSM). When ESM is required, it is explicitly stated.

Review of the DHCP Protocol

The DHCP protocol is used by a DHCP server to dynamically assign IP addresses and other optional configuration parameters on request of DHCP clients. These parameters are leased by the DHCP server for a duration specified by the lease time.

The DHCP lease process is outlined in Figure 162.

Figure 162 **DHCP Lease Process**

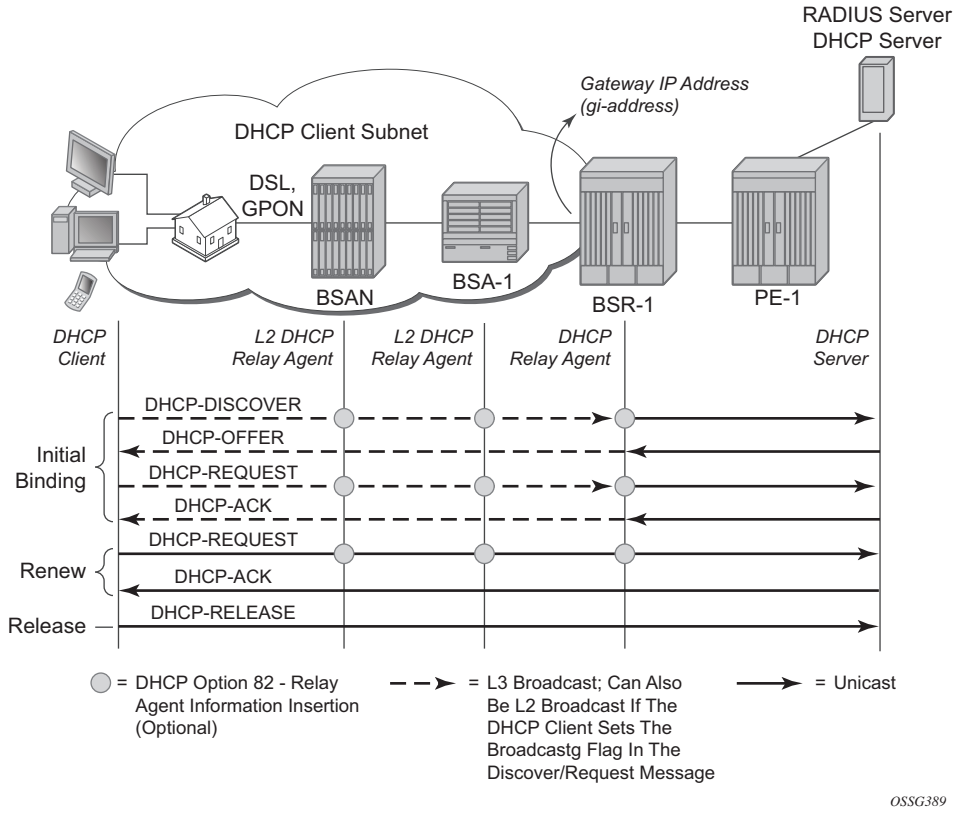

When a DHCP client boots, a DHCP discover message is broadcast on the local subnet (dest-ip = 255.255.255.255).

A DHCP server in the local subnet responds with a unicast DHCP offer message containing the *your ip address* field as well as other configuration parameters in the option fields (such as subnet mask, default gateway, DNS server IP addresses, lease time, etc.).

The DHCP client responds with a DHCP request message to accept the parameters specified in the DHCP offer. The DHCP request is also broadcast on the local subnet.

The DHCP server acknowledges the DHCP request with a unicast DHCP ack message.

When the DHCP client receives a DHCP ack from the server, it is said to be in the bound state.

When half of the lease time has expired, the DHCP client tries to renew the lease. It will send a unicast DHCP request message to the DHCP server. The DHCP server will reply to the request with a unicast DHCP ack to the client.

If the renew failed, a rebind is attempted by default at 7/8 of the lease time. It will send a broadcast DHCP request message.

Before disconnecting from the local subnet, a DHCP client may return its lease by sending a DHCP release message to the DHCP server.

In case there is no DHCP server in the subnet of the DHCP client, a DHCP relay agent is needed to forward the broadcast DHCP discover/request messages on behalf of the DHCP client to a DHCP server located on a different subnet. The DHCP relay agent will add the gateway IP address field to the messages and send them as unicast to the DHCP server IP address. The DHCP server in this case will respond by unicast to the DHCP relay agent. The DHCP relay agent forwards the DHCP server messages as broadcast on the DHCP client subnet.

Configuration

DHCP Snooping

DHCP client originated messages (discover, request, release) must be snooped (intercepted and sent to the control plane for further processing) to allow for DHCP Option 82 insertion, authentication through local user database (LUDB), AAA/ RADIUS or AAA/Diameter, and releasing the DHCP host session state.

For Bridged CO, DHCP snooping must be enabled explicitly on the subscriber SAP:

Bridged CO @ BSA-1:

```
configure
    service
        vpls 1
            --- snipped ---
            sap 1/1/3:1 split-horizon-group "rshg-1" create
                description "sub-1"
                dhcp
```

```
                    snoop
                    no shutdown
                exit
            exit
        exit
    exit
exit
```

DHCP server originated messages (offer, ack, nak, etc.) must be snooped to allow for DHCP Option 82 removal, lease state population and/or ESM functions.

For Bridged CO, DHCP snooping must be enabled explicitly on all SDPs and/or SAPs that provide connectivity to the DHCP server:

Bridged CO @ BSA-1:

```
configure
    service
        vpls 1
            --- snipped ---
            spoke-sdp 1:1 create
                dhcp
                    snoop
                exit
                no shutdown
            exit
        exit
    exit
exit
```

For Routed CO, DHCP snooping is implicitly enabled by configuring a DHCP relay agent (DHCP Relay Agent): All DHCP messages received on a routed network interface will be snooped (for example, intercepted and sent to the control plane for further processing).

DHCP Relay Agent

For Bridged CO, the DHCP relay agent function is configured inthe IP edge (BSR), at the regular interface):

Bridged CO @ BSR-1:

```
configure
    service
        vprn 1
            --- snipped ---
            interface "int-BSA1-p2mp-1" create
                description "Bridged CO"
                address 10.1.0.254/16
                dhcp
```

```
                    server 172.16.0.1
                    trusted
                    gi-address 10.1.0.254
                    no shutdown
                exit
                --- snipped ---
                ip-mtu 1500
                spoke-sdp 1:1 create
                    no shutdown
                exit
            exit
        exit
    exit
exit
```

For Routed CO, the DHCP relay agent function must be configured at BSR-1 group-interface level where the DHCP host will be instantiated:

Routed CO @ BSR-1:

```
configure
    service
        vprn 1
            --- snipped ---
            subscriber-interface "sub-int-1" create
                description "Routed CO"
                address 10.2.0.254/16
                group-interface "group-int-1" create
                    --- snipped ---
                    dhcp
                        server 172.16.0.1
                        trusted
                        --- snipped ---
                        gi-address 10.2.0.254
                        no shutdown
                    exit
                exit
            exit
        exit
    exit
exit
```

The server IP address should point to the DHCP server and must be reachable in the same routing instance as where the (subscriber-)interface is defined.

The **trusted** command makes the interface a trusted interface to allow Option 82 insertion by a Layer 2 DHCP relay agent (see DHCP Options (Relay Agent Information)).

The gi-address must be a local configured IP address on the (subscriber-)interface. The DHCP messages relayed to the DHCP server will have the outgoing interface IP address as source IP address. To change the default behavior and use the configured gi-address as source IP address, specify the optional **src-ip-addr** flag:

CLI Syntax: `gi-address 10.2.0.254 src-ip-addr`

A Layer 2 DHCP relay agent (such as BSAN or BSA) can add DHCP Option 82 information and leave the gi-address field to 0.0.0.0. The gi-address is the gateway IP address, filled in by the DHCP relay agent. An incoming DHCP discover with Option 82 present and gi-address field = 0.0.0.0 will be dropped by the DHCP relay agent according the RFC. The Rx Untrusted Packets and client Packets Discarded counters are increased in the DHCP statistics.

Output from DHCP debug log on BSR-1:

```
2 2016/02/24 09:06:12.92 CET MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 4 (group-int-1),
   DROPPED DHCP Boot Request on Interface group-int-1 (1/1/3:1) Port 67
   Problem: message is received from an untrusted client
```

Therefore, the DHCP relay agent should be configured as trusted to allow DHCP Option 82 insertion by a Layer 2 DHCP relay agent.

DHCP Options (Relay Agent Information)

In Bridged CO, when DHCP snooping is enabled on a VPLS SAP, DHCP Option 82 relay agent information can be altered or added on an incoming DHCP discover/ request. This is sometimes referred to as a Layer 2 DHCP relay agent function.

In Routed CO, a DHCP relay agent can alter or add the DHCP Option 82 relay agent information on an incoming DHCP discover/request.

Supported DHCP Option 82 sub-options and their format are listed in Table 33:

Table 33 **Supported DHCP Option 82 Sub-Options**

Option 82 Sub-Option	Format	Example
Opt82 [1] Circuit ID (Routed CO)	ifindex — 32 bit virtual router ID followed by a 32 bit ifindex in hex	00 00 00 02 00 00 00 04
	sap-id [sap id in ascii]	1/1/3:1
	ascii3-tuple [system-name|service-id|group-interface|sap-id]	
	vlan-ascii-tuple [system-name|service-id|group-interface|dot1p|vlan-id]	“BSR-1|1|group-int-1|0|1”

Table 33 **Supported DHCP Option 82 Sub-Options (Continued)**

Option 82 Sub-Option	Format	Example
Opt82 [1] Circuit ID (Bridged CO)	ascii-tuple [system-name|service-id|sap-id]	“BSA-1|1|1/1/2:1”
	vlan-ascii-tuple [system-name|service-id|sap-id |dot1p|vlan-id]	“BSA-1|1|1/1/2:1|0|1”
Opt82 [2] Remote ID (Bridged and Routed CO)	MAC [client hw address in hex]	fe fd 00 02 45 00
	string (max. 32 chars)	“Opt-82 [2] – Remote ID”
Opt82 [9] Vendor Specific (Bridged and Routed CO)	[1] system-id [hostname in ascii]	“BSA-1” or “BSR-1”
	[2] client-mac-address [client hw address in hex]	fe fd 00 02 45 00
	[3] service-id	1
	[4] sap-id [sap id in ascii]	“1/1/2:1”
	[5] string (max. 32 chars)	“Opt-82 [9] [5] – string”
Opt82 [9] Vendor Specific (Routed CO)	[13] pool-name [dhcp pool name from Radius/ Local User DB in ascii	“dhcp-pool-1”

Note: The application for Option 82 Circuit-ID format vlan-ascii-tuple is to preserve the Dot1p marking of DHCP packets in the downstream direction (DHCP server to client). The dot1p value of the incoming DHCP discover/request is recorded as part of the Option 82 Circuit ID. The outgoing DHCP offer/ack packets are marked with the Dot1p value found as part of the Circuit ID echoed by the DHCP server.

Following actions can be taken on incoming DHCP discover/request:

- replace
- drop
- keep (default)

Replace

At ingress:

If present, remove all the Option 82 information from the incoming DHCP discover/ request. Insert the configured DHCP options before forwarding to the DHCP relay agent or DHCP server

At egress:

Remove all Option 82 information from the incoming DHCP offer/ack before forwarding to the client.

Bridged CO @ BSA-1:

```
configure
    service
        vpls 1
            --- snipped ---
            sap 1/1/3:1 split-horizon-group "rshg-1" create
                description "sub-1"
                dhcp
                    snoop
                    option
                        action replace
                        remote-id string "Opt-82 [2] - Remote ID"
                        vendor-specific-option
                            system-id
                            client-mac-address
                            service-id
                            sap-id
                            string "Opt-82 [9][5] - Vendor ID"
                        exit
                    exit
                    no shutdown
                exit
            exit
         exit
     exit
exit
```

Routed CO @ BSR-1:

```
configure
    service
        vprn 1
            --- snipped ---
            subscriber-interface "sub-int-1" create
                description "Routed CO"
                address 10.2.0.254/16
                group-interface "group-int-1" create
                    --- snipped ---
                    dhcp
                        option
                            action replace
                            circuit-id
                            remote-id string "Opt-82 [2] Remote-ID"
                            vendor-specific-option
                                system-id
                                client-mac-address
                                pool-name
                                service-id
                                sap-id
                                string "Opt-82 [9][5] string"
                            exit
```

```
                        exit
                        server 172.16.0.1
                        trusted
                        --- snipped ---
                        gi-address 10.2.0.254
                        no shutdown
                    exit
                exit
            exit
         exit
     exit
exit
```

Drop

Drop all incoming DHCP discover/request with Option 82 information present.

Incoming DHCP discover/request without Option 82 information will be forwarded to (Bridged CO) or processed by (Routed CO) the DHCP relay agent as is, ignoring the configured options.

Bridged CO @ BSA-1:

```
configure
    service
        vpls 1
        --- snipped ---
                sap 1/1/3:1 split-horizon-group "rshg-1" create
                description "sub-1"
                dhcp
                    snoop
                    option
                        action drop
                    exit
                exit
            exit
        exit
    exit
exit
```

Routed CO @ BSR-1:

```
configure
    service
        vprn 1
            --- snipped ---
            subscriber-interface "sub-int-1" create
                description "Routed CO"
                address 10.2.0.254/16
                group-interface "group-int-1" create
                    --- snipped ---
                    dhcp
                        option
```

```
                            action drop
                        exit
                        server 172.16.0.1
                        trusted
                        gi-address 10.2.0.254
                        no shutdown
                    exit
                exit
            exit
        exit
    exit
exit
```

Output from the DHCP debug log:

Bridged CO @ BSA-1:

```
2 2016/02/24 09:08:34.10 CET MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: Dropped DHCP Packet
   VPLS 1, SAP 1/1/3:1
   Problem: port config doesn't allow BOOTP/DHCP packets with option 82
```

Routed CO @ BSR-1:

```
4 2016/02/24 09:08:58.17 CET MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 4 (group-int-1),
   DROPPED DHCP Boot Request on Interface group-int-1 (1/1/3:1) Port 67
   Problem: action drop is configured and packet contains option 82
```

The Clients Packets Dropped counter is increased in the DHCP statistics:

Bridged CO @ BSA-1:

```
*A:BSA-1# show service id 1 dhcp statistics
=====================================================
DHCP Statistics, service 1
=====================================================
Client Packets Snooped               : 4
Client Packets Forwarded             : 4
Client Packets Dropped               : 3
Client Packets Proxied (RADIUS)      : 0
Client Packets Proxied (Diameter)    : 0
Client Packets Proxied (User-Db)     : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Snooped               : 3
Server Packets Forwarded             : 3
Server Packets Dropped               : 0
DHCP RELEASEs Spoofed                : 0
DHCP FORCERENEWs Spoofed             : 0
=====================================================
*A:BSA-1#
```

Routed CO @ BSR-1:

```
*A:BSR-1# show service id 1 dhcp statistics
===============================================================================
DHCP Global Statistics, service 1
===============================================================================
Rx Packets                          : 326
Tx Packets                          : 318
Rx Malformed Packets                : 0
Rx Untrusted Packets                : 1
Client Packets Discarded            : 7
Client Packets Relayed              : 80
Client Packets Snooped              : 82
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              : 80
Server Packets Snooped              : 76
DHCP RELEASEs Spoofed               : 0
DHCP FORCERENEWs Spoofed            : 0
===============================================================================
*A:BSR-1#
```

Keep (default)

At ingress — Incoming DHCP discover/request without Option 82 information will be forwarded to (Bridged CO) or processed by (Routed CO) the DHCP relay agent as is, ignoring any configured option.

At ingress for incoming DHCP discover/request with Option 82 information present —Configured vendor specific options will be merged with the existing Option 82 information before sending to (Routed CO) or processing by (Routed CO) the DHCP relay agent. Configured Circuit ID and Remote ID options will be ignored.

At egress — Remove Option 82 vendor specific information from the incoming DHCP offer/ack before forwarding to the client. Other existing DHCP Option 82 information is retained.

Bridged CO @ BSA-1:

```
configure
    service
        vpls 1
            --- snipped ---
            sap 1/1/3:1 split-horizon-group "rshg-1" create
                description "sub-1"
                dhcp
                    snoop
                    option
                        action keep
```

```
                exit
            exit
        exit
    exit
exit
```

Routed CO @ BSR-1:

```
configure
    service
        vprn 1
            --- snipped ---
            subscriber-interface "sub-int-1" create
                description "Routed CO"
                address 10.2.0.254/16
                group-interface "group-int-1" create
                    --- snipped ---
                    dhcp
                        option
                            action keep
                        exit
                        server 172.16.0.1
                        trusted
                        gi-address 10.2.0.254
                        no shutdown
                    exit
                exit
            exit
        exit
    exit
exit
```

DHCP Lease State

The DHCP lease state table is an internal database that keeps track of the DHCP host states. The DHCP lease state table enables subscriber management functions (per-subscriber QoS and accounting) and security functions (dynamic anti-spoof filtering) on the DHCP host.

The DHCP lease information for a specific host is extracted from the DHCP ack message.
Table 34 displays some typical information stored in the DHCP lease state. The table does not display all information: additional data is added for managed SAPs, DHCPv6, etc.

Table 34 **Information in DHCP Lease State**

Parameter	Comment
Service ID	Service where the DHCP host is connected
IP Address	IP address of the DHCP host
Client HW Address	Ethenet MAC address of the DHCP host
Subscriber-interface (Routed CO only)	Subscriber interface name where the DHCP host is instantiated
Group-interface (Routed CO only)	Group interface name where the DHCP host is instantiated
SAP	SAP where the DHCP hosts is connected
Up Time	The DHCPv4 host uptime
Remaining Lease Time	The time remaining before the lease expires
Remaining SessionTime	The time remaining before the DHCPv4 host is deleted from the system (updated each time a DHCP renew/rebind occurs)
Persistence Key	Lookup key for this host in the persistency file (see further)
Sub-Ident	ESM: Subscriber ID of the DHCP host
Sub-Profile-String	ESM: Subscriber profile string of the DHCP host
SLA-Profile-String	ESM: SLA profile string of the DHCP host
App-Profile-String	ESM: Application profile string of the DHCP host
Lease ANCP-String	ESM: ANCP string for this DHCP host
Lease Int Dest Id	ESM: Internal destination ID for this DHCP host
Category-Map-Name	ESM: Volume and Time based accounting
Lease Info origin	ESM: Origin for the IP configuration for this host (None, DHCP, RADIUS, etc.)
Ip-Netmask	The IP netmask for this DHCP host
Broadcast-Ip-Addr	The broadcast IP address for this host
Default-Router	The default gateway for this host
Primary-Dns	The primary DNS server for this host
Secondary-Dns	The secondary DNS server for this host
Primary-Nbns	The primary NetBIOS name server for this host
Secondary-Nbns	The secondary NetBIOS name server for this host

Table 34 **Information in DHCP Lease State (Continued)**

Parameter	Comment
ServerLeaseStart	Time and date that the lease for this host started (first DHCP ack received)
ServerLastRenew	Time and date that the lease for this host was last renewed
ServerLeaseEnd	Time and date that the lease for this host will expire
Session-Timeout	The DHCPv4 is deleted when its uptime reaches this value
IPoE|PPP session	Indication if this lease belongs to an IPoE or PPP session, or to no session
Lease-Time	The lease time specified by the DHCPv4 server
DHCP Server Addr	IP address of the DHCP server that allocated the lease for this host
Circuit Id	DHCP Relay Agent information Option 82 Circuit ID content
Remote Id	DHCP Relay Agent information Option 82 Remote ID content
RADIUS User-Name	ESM: Username used in the RADIUS authentication access request

For Bridged CO, populating the DHCP lease state table must be explicitly enabled through configuration. The number of leases allowed on the VPLS SAP must be specified. When omitted, a single DHCP host is allowed per SAP.

Bridged CO @ BSA-1:

```
configure
    service
        vpls 1
            --- snipped ---
            sap 1/1/3:1 split-horizon-group "rshg-1" create
                description "sub-1"
                dhcp
                    snoop
                    lease-populate 10
                    no shutdown
                exit
            exit
         exit
     exit
exit
```

For Routed CO, DHCP lease state table population is enabled by default on a group interface with DHCP configured as **no shutdown** The number of leases allowed on each SAP of the group-interface must be configured (by default a single DHCP host is allowed on each SAP):

Routed CO @ BSR-1:

```
configure
    service
        vprn 1
            --- snipped ---
            subscriber-interface "sub-int-1" create
                description "Routed CO"
                address 10.2.0.254/16
                group-interface "group-int-1" create
                    dhcp
                        server 172.16.0.1
                        trusted
                        lease-populate 10
                        gi-address 10.2.0.254
                        no shutdown
                    exit
                exit
            exit
         exit
     exit
exit
```

To check the DHCP lease state for a particular service, use the **show service id** *service-id* **dhcp lease-state** command. Detailed output as well as additional output filtering is available:

Bridged CO:

```
*A:BSA-1# show service id 1 dhcp lease-state ?
  - 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]
```

Routed CO:

```
A:BSR-1# show service id 1 dhcp lease-state ?
  - 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]
```

Bridged CO @ BSA-1:

```
*A:BSA-1# show service id 1 dhcp lease-state mac 00:00:00:11:11:11 detail
===============================================================================
DHCP lease states for service 1
===============================================================================
Service ID           : 1
IP Address           : 10.1.0.100
Client HW Address    : 00:00:00:11:11:11
SAP                  : 1/1/3:1
Termination Type     : local
Up Time              : 0d 00:00:40
Remaining Lease Time : 0d 11:59:21
Remaining SessionTime: N/A
```

```
Persistence Key      : N/A

Sub-Ident            : "sub-11"
Sub-Profile-String   : "sub-prof-1"
SLA-Profile-String   : "sla-prof-1"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : ""
Category-Map-Name    : ""

Lease Info origin    : DHCP

Ip-Netmask           : 255.255.0.0
Broadcast-Ip-Addr    : 10.1.255.255
Default-Router       : 10.1.0.254
Primary-Dns          : N/A
Secondary-Dns        : N/A
Primary-Nbns         : N/A
Secondary-Nbns       : N/A

ServerLeaseStart     : 02/24/2016 09:11:08
ServerLastRenew      : 02/24/2016 09:11:08
ServerLeaseEnd       : 02/24/2016 21:11:08
Session-Timeout      : N/A
IPoE|PPP session     : No
Lease-Time           : 0d 12:00:00
DHCP Server Addr     : 172.16.0.1

Relay Agent Information
  Circuit Id         : 11
Radius User-Name     : "00:00:00:11:11:11"
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
*A:BSA-1#
```

Routed CO:

```
*A:BSR-1# show service id 1 dhcp lease-state mac 00:00:00:33:33:31 detail
===============================================================================
DHCP lease states for service 1
===============================================================================
Service ID           : 1
IP Address           : 10.2.0.106
Client HW Address    : 00:00:00:33:33:31
Subscriber-interface : sub-int-1
Group-interface      : group-int-1
SAP                  : 1/1/3:1
Up Time              : 0d 00:01:11
Remaining Lease Time : 0d 00:08:49
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "sub-31"
Sub-Profile-String   : "sub-prof-2"
SLA-Profile-String   : "sla-prof-2"
App-Profile-String   : ""
Lease ANCP-String    : ""
```

```
Lease Int Dest Id    : ""
Category-Map-Name    : ""

Lease Info origin    : DHCP

Ip-Netmask           : 255.255.0.0
Broadcast-Ip-Addr    : 10.2.255.255
Default-Router       : 10.2.0.254
Primary-Dns          : N/A
Secondary-Dns        : N/A
Primary-Nbns         : N/A
Secondary-Nbns       : N/A

ServerLeaseStart     : 02/24/2016 09:11:30
ServerLastRenew      : 02/24/2016 09:11:30
ServerLeaseEnd       : 02/24/2016 09:21:30
Session-Timeout      : N/A
Lease-Time           : 0d 00:10:00
DHCP Server Addr     : 172.16.0.1

Relay Agent Information
  Circuit Id         : 55
Radius User-Name     : "00:00:00:33:33:31"
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
*A:BSR-1#
```

DHCP Host Session: Set-up, Operation and Release

Snooping the DHCP communication between a DHCP client and a DHCP relay agent/server facilitates the DHCP host instantiation: Upon the reception of a DHCP ack message from the server, the DHCP lease state table is populated. With ESM enabled, a DHCP host is also instantiated. The DHCP host will appear in the subscriber-host table for the service with origin set to DHCP.

Bridged CO:

```
*A:BSA-1# show service id 1 subscriber-hosts
===============================================================
Subscriber Host table
===============================================================
Sap                   Subscriber
  IP Address
    MAC Address           PPPoE-SID Origin
---------------------------------------------------------------
1/1/3:1               sub-11
  10.1.0.100
    00:00:00:11:11:11    N/A       DHCP
---------------------------------------------------------------
Number of subscriber hosts : 1
===============================================================
*A:BSA-1#
```

Routed CO:

```
*A:BSR-1# show service id 1 subscriber-hosts
===============================================================
Subscriber Host table
===============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID Origin       Fwding State
---------------------------------------------------------------
1/1/3:1                sub-31
  10.2.0.106
    00:00:00:33:33:31    N/A       DHCP         Fwding
---------------------------------------------------------------
Number of subscriber hosts : 1
===============================================================
*A:BSR-1#
```

If ESM is enabled, the subscriber-host will also appear in the active subscriber table:

Routed CO:

```
*A:BSR-1# show service active-subscribers
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber sub-31 (sub-profile-1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:1/1/3:1 - sla:sla-profile-1
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.2.0.106
                00:00:00:33:33:31 N/A       DHCP
-------------------------------------------------------------------------------
Number of active subscribers : 1
-------------------------------------------------------------------------------
*A:BSR-1#
```

Troubleshooting the DHCP session set-up is done with DHCP debugging:

Bridged CO:

```
*A:BSA-1# debug service id 1 dhcp ?
  - dhcp
  - no dhcp
 [no] detail-level    - Configure the DHCP tracing detail level
 [no] mac             - Show DHCP packets for a particular MAC address
 [no] mode            - Configure the DHCP tracing mode
 [no] sap             - Show DHCP packets for a particular SAP
 [no] sdp             - Show DHCP packets for a particular SDP
*A:BSA-1#
```

Routed CO:

```
*A:BSR-1# debug router 1 ip dhcp ?
  - dhcp [interface <ip-int-name>]
  - dhcp mac <ieee-address>
  - dhcp sap <sap-id>
  - no dhcp [interface <ip-int-name>]
  - no dhcp mac <ieee-address>
  - no dhcp sap <sap-id>
    --- snipped ---
*A:BSR-1#
```

For example:

Bridged CO:

```
*A:BSA-1# show debug
debug
    service
        id 1
            dhcp
                mode egr-ingr-and-dropped
                detail-level medium
            exit
        exit
    exit
exit
*A:BSA-1#
```

Routed CO:

```
*A:BSR-1# show debug
debug
    router "1"
        ip
            dhcp
                detail-level medium
                mode egr-ingr-and-dropped
            exit
        exit
    exit
exit
*A:BSR-1#
```

The example above will log all DHCP packets on the service. When thousands of DHCP hosts are active, more granular filtering is required: for example look only to dropped packets or look only to packets from a particular MAC address.

To display the debugging information, a dedicated log should be created:

```
*A:BSA-1# configure log
*A:BSA-1>config>log# info
----------------------------------------------
        log-id 1
            description "Send debug log to the current telnet/ssh session"
```

```
            from debug-trace
            to session
            no shutdown
         exit
-----------------------------------------------
*A:BSA-1>config>log#
```

The following shows sample DHCP debug log output (detail-level medium):

Bridged CO:

```
20 2016/02/24 09:15:58.20 CET MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: TX DHCP Packet
   VPLS 1, SAP 1/1/3:1
   BootReply to UDP port 68
   ciaddr: 0.0.0.0           yiaddr: 10.1.0.100
   siaddr: 172.16.0.1        giaddr: 10.1.0.254
   chaddr: 00:00:00:11:11:11    xid: 0x9
   DHCP options:
   [82] Relay agent information: len = 4
      [1] Circuit-id: 11
   [53] Message type: Ack
   [54] DHCP server addr: 172.16.0.1
   [51] Lease time: 43200
   [1] Subnet mask: 255.255.0.0
   [3] Router: 10.1.0.254
   [28] Broadcast addr: 10.1.255.255
   [255] End
"
```

During the lifetime of a DHCP host, the DHCP lease state is updated in the system: for example, the remaining lifetime after a DHCP renew. To check lease details from the DHCP host during its lifetime, consult the DHCP lease state details:

```
*A:BSA-1# show service id 1 dhcp lease-state detail
===============================================================================
DHCP lease states for service 1
===============================================================================
Service ID           : 1
IP Address           : 10.1.0.100
Client HW Address    : 00:00:00:11:11:11
--- snipped ---
*A:BSA-1#
```

If the remaining lifetime timer expires before the DHCP session is renewed or rebound, the DHCP lease state is cleared. If ESM is enabled, the DHCP host is removed from the system.

A DHCP host can be manually deleted from the system using following clear command:

```
*A:BSA-1# clear service id 1 dhcp lease-state ?
  - 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 port <port-id> [no-dhcp-release]
  - lease-state sap <sap-id> [no-dhcp-release]
  - lease-state sdp <sdp-id:vc-id> [no-dhcp-release]

*A:BSA-1# clear service id 1 dhcp lease-state ip-address 10.1.0.100
```

The DHCP lease state is deleted with all related state (such as, anti-spoof filter, ARP table entry). If ESM is enabled, the DHCP host is removed from the system. Optionally, a DHCP release is sent to the DHCP server to notify that the IP address can be released. This is reflected in the DHCP statistics in the DHCP RELEASEs Spoofed counter. Use the **no-dhcp-release** flag in the clear command if no DHCP release is to be sent when issuing the **clear** command.

To display a summary overview of the DHCP configuration on a particular service:

Bridged CO:

```
*A:BSA-1# show service id 1 dhcp summary
===============================================================================
DHCP Summary, service 1
===============================================================================
Sap/Sdp                 Snoop  Used/       Arp Reply   Info     Admin
                               Provided    Agent       Option   State
-------------------------------------------------------------------------------
sap:1/1/3:1             Yes    0/10        Yes         Keep     Up
sap:1/1/3:2             Yes    0/10        Yes         Keep     Up
sdp:1:1                 Yes    N/A         N/A         N/A      N/A
-------------------------------------------------------------------------------
Number of Entries : 3
-------------------------------------------------------------------------------
===============================================================================
*A:BSA-1#
```

Routed CO:

```
*A:BSR-1# show service id 1 dhcp summary
===============================================================================
DHCP Summary, service 1
===============================================================================
Interface Name                  Arp      Leases Per Interface/ Info    Admin
  SapId/Sdp                     Populate Per Sap Limit         Option  State
-------------------------------------------------------------------------------
group-int-1                     Yes      1/10                  Keep    Up
int-BSA1-p2mp-1                 No       0/0                   Keep    Up
-------------------------------------------------------------------------------
Interfaces: 2
===============================================================================
*A:BSR-1#
```

The Leases Per Interface/Per Sap Limit field indicates the number of active versus the number of allowed DHCP leases on the SAP, SDP or interface.

To check the DHCP statistics, use the following command:

Bridged CO:

```
*A:BSA-1# show service id 1 dhcp statistics
=====================================================
DHCP Statistics, service 1
=====================================================
Client Packets Snooped               : 10
Client Packets Forwarded             : 9
Client Packets Dropped               : 1
Client Packets Proxied (RADIUS)      : 0
Client Packets Proxied (Diameter)    : 0
Client Packets Proxied (User-Db)     : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Snooped               : 7
Server Packets Forwarded             : 7
Server Packets Dropped               : 0
DHCP RELEASEs Spoofed                : 1
DHCP FORCERENEWs Spoofed             : 0
=====================================================
*A:BSA-1#
```

Routed CO:

```
*A:BSR-1# show service id 1 dhcp statistics
=====================================================================
DHCP Global Statistics, service 1
=====================================================================
Rx Packets                           : 341
Tx Packets                           : 332
Rx Malformed Packets                 : 0
Rx Untrusted Packets                 : 1
Client Packets Discarded             : 8
Client Packets Relayed               : 86
Client Packets Snooped               : 83
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               : 86
Server Packets Snooped               : 77
DHCP RELEASEs Spoofed                : 0
DHCP FORCERENEWs Spoofed             : 0
=====================================================================
*A:BSR-1#
```

Note: Additional filtering can be done to retrieve DHCP statistics per SAP, SDP or interface.

To clear the DHCP statistics:

Bridged CO:

```
*A:BSA-1# clear service id 1 dhcp statistics ?
  - statistics [sap <sap-id> | sdp <sdp-id:vc-id> | interface <ip-int-name|ip-
address>]
```

Routed CO:

```
*A:BSR-1# clear router 1 dhcp statistics ?
  - statistics [<ip-int-name|ip-address>]
 <ip-int-name|ip-ad*> : ip-int-name    - 32 chars max
                        ip-address     - a.b.c.d
```

DHCP Hosts Advanced Topics

High Availability

The DHCP lease state supports High Availability (HA): the lease state table is synchronized to the standby CPM. When the active CPM fails, all DHCP hosts stay active without service interruption.

DHCP Lease State Persistency

A DHCP session does not have a keep-alive mechanism to detect unavailability. A new DHCP session set up is only attempted after expiration of the DHCP lease time. A node reboot causing the loss of DHCP lease state and the corresponding anti-spoof filters could therefore result in unacceptable long service outages.

The DHCP lease state can be made persistent across node reboots: DHCP lease state is restored from a persistency file stored on the compact flash file system. As a result, DHCP sessions will only loose connectivity during the time of reboot without being completely disconnected.

To activate the DHCP lease state persistency:

```
configure
    system
        persistence
            subscriber-mgmt
                description "DCHP lease state persistency"
                location cf3:
```

```
            exit
        exit
    exit
exit
```

A dedicated persistency file will be created on the specified compact flash file system. The file is initialized to store the maximum number of allowed hosts; its size is fixed to avoid file system space problems during operations.

```
*A:BSA-1# file dir cf3:\sub*
Volume in drive cf3 on slot A is SROS VM.
Volume in drive cf3 on slot A is formatted as FAT32
Directory of cf3:
02/25/2016  04:44a           536871424 submgmt.011
02/25/2016  04:44a            12583424 submgmt.i11
               2 File(s)              549454848 bytes.
               0 Dir(s)               330903552 bytes free.
*A:BSA-1#
```

Each time a DHCP ack is received from the DHCP server, the persistency file is updated together with the lease state. If the file update fails, an event is generated to indicate that persistency can not be guaranteed.

The content of the persistency file may vary between different SR-OS software releases. When upgrading, the persistency file is automatically upgraded to the new format. To downgrade the persistency file to a lower SR-OS release version, use the following command:

```
*A:BSA-1# tools perform persistence downgrade ?
  - downgrade target-version <target> [reboot]
 <target>             : the version you want to downgrade to
                        submgt
                          13.0 (current)  - cf3:\submgmt.011
                          12.0            - cf3:\submgmt.010
                          11.0            - cf3:\submgmt.009
                          10.0            - cf3:\submgmt.008
                          9.0             - cf3:\submgmt.007
                          8.0             - cf3:\submgmt.006
                          7.0             - cf3:\submgmt.005
                          6.0             - cf3:\submgmt.004
                          5.0             - cf3:\submgmt.003
                          4.0             - cf3:\submgmt.pst
 <reboot>             : reboot system after successful conversion
```

The content of the persistency file can be looked at using the following command:

```
*A:BSA-1# show service id 1 dhcp lease-state mac 00:00:00:11:11:11 detail
===============================================================================
DHCP lease states for service 1
===============================================================================
Service ID           : 1
IP Address           : 10.1.0.100
Client HW Address    : 00:00:00:11:11:11
SAP                  : 1/1/3:1
```

```
Termination Type     : local
Up Time              : 0d 00:00:15
Remaining Lease Time : 0d 11:59:45
Remaining SessionTime: N/A
Persistence Key      : 0x00000000

--- snipped ---

Relay Agent Information
  Circuit Id         : 11
Radius User-Name     : "00:00:00:11:11:11"
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
*A:BSA-1#


*A:BSA-1# tools dump persistence submgt record 0x00000000
---------------------------------------
Persistence Record
---------------------------------------
Client      : submgt
Persist-Key : 0x00000000
Filename    : cf3:\submgmt.011
Entries     : Index  FedHandle  Last Update               Action Valid
              000064 0x00000000 2016/02/24 08:21:13 (UTC) ADD    Yes
Data        : 300 bytes
 Host Type      : DHCP lease state
 Service ID     : 1
 SAP ID         : 1/1/3:1
 NH MAC         : 00:00:00:11:11:11
 Srvr Lse Start : 2016/02/24 08:21:13 (UTC)
 IP             : 10.1.0.100
 CHADDR         : 00:00:00:11:11:11
 Srvr Last Renew: 2016/02/24 08:21:13 (UTC)
 Srvr Lse End   : 2016/02/24 20:21:13 (UTC)
 Srvr Addr      : 172.16.0.1
 Option82       : 4 bytes
 RADIUS Fallback: NO
 Acct-Sess-Id   : 02D9FF0000000856CD67F9
 Multi-Sess-Id  : 02D9FF0000000956CD67F9
 Class Attr     : 0 bytes
 User-Name      : "00:00:00:11:11:11"
 Address Origin : DHCP
 host is authenticated by radius: true
 Subscriber-Id  : "sub-11"
 Sub-Profile-Str: "sub-prof-1"
 SLA-Profile-Str: "sla-prof-1"
 Framed IP Netmask: 255.255.0.0
 Broadcast IP Address: 10.1.255.255
 Default Router : 10.1.0.254
 Lease-Time     : 43200
*A:BSA-1#
```

Limiting the Number of DHCP Hosts

Lease populate limit

The maximum number of DHCP lease state entries for a VPLS SAP, for an IES/VPRN interface or for each SAP on an IES/VPRN group-interface is defined when enabling the lease-populate. When omitted, a single DHCP host is allowed:

```
configure
    service
        vpls 1
            --- snipped ---
            sap 1/1/3:2 split-horizon-group "rshg-1" create
                description "sub-1"
                dhcp
                    snoop
                    lease-populate 1
                    no shutdown
                exit
            exit
        exit
    exit
exit
```

When trying to instantiate a new DHCP host while the configured number of leases is reached, the DHCP ack is dropped (DHCP debug log output):

```
44 2016/02/24 09:24:30.67 CET MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: Dropped DHCP Packet
   VPLS 1, spoke-sdp 1:1
   --- snipped ---
   Problem: lease-populate limit (1) exceeded on SAP 1/1/3:2
```

The following event is generated (log-id 99):

```
71 2016/02/
24 09:24:30.67 CET WARNING: DHCP #2002 Base Maximum number of lease states *
"Lease state for (CiAddr = 10.1.0.103, ChAddr = 00:00:00:22:22:22, leaseTime = 43200
) was not stored because the number of DHCP lease states on SAP 1/1/
3:2 in service 1 has reached its upper limit"
```

With ESM enabled, the following additional limits apply:

- sla-profile host-limits
- multi-sub-sap limit

SLA-profile host limits

TheSLA-profile contains host limits defining the maximum number of dynamic subscriber hosts per subscriber for this sla-profile. Static hosts are not counted in the host-limits.

```
*A:BSA-1>config>subscr-mgmt>sla-prof# host-limits ?
  - host-limits
  - no host-limits
 [no] ipv4-arp        - Maximum number of IPv4 ARP hosts
 [no] ipv4-dhcp       - Maximum number of IPv4 DHCP hosts
 [no] ipv4-overall    - Maximum number of IPv4 hosts
 [no] ipv4-ppp        - Maximum number of IPv4 PPP hosts
 [no] ipv6-overall    - Maximum number of IPv6 hosts
 [no] ipv6-pd-ipoe-d* - Maximum number of IPv6-PD IPOE DHCP hosts
 [no] ipv6-pd-overall - Maximum number of IPv6-PD hosts
 [no] ipv6-pd-ppp-dh* - Maximum number of IPv6-PD PPP DHCP hosts
 [no] ipv6-wan-ipoe-* - Maximum number of IPv6-Wan IPOE DHCP hosts
 [no] ipv6-wan-ipoe-* - Maximum number of IPv6-Wan IPOE SLAAC hosts
 [no] ipv6-wan-overa* - Maximum number of IPv6-Wan hosts
 [no] ipv6-wan-ppp-d* - Maximum number of IPv6-Wan PPP DHCP hosts
 [no] ipv6-wan-ppp-s* - Maximum number of IPv6-Wan PPP SLAAC hosts
 [no] lac-overall     - Maximum number of L2TP LAC hosts
 [no] overall         - Maximum number of hosts
 [no] remove-oldest   - Remove oldest
```

Optionally the remove-oldest command can be used. In that case, the new host is accepted and the DHCP lease state for the oldest host (with the least remaining lease time) is cleared. A DHCP release message is sent to the DHCP server.

The following example limits the amount of ipv4-dhcp hosts.

```
configure
    subscriber-mgmt
        sla-profile "sla-profile-2" create
            host-limits
                ipv4-dhcp 1
            exit
        exit
    exit
exit
```

If the configured host-limit is reached for a subscriber, access is denied for a new host, an event is generated (log-id 99) and the corresponding DHCP ack message is dropped:

```
80 2016/02/24 14:18:46.46 CET WARNING: DHCP #2005 Base Lease State Population Error
"Lease state table population error on SAP 1/1/3:2 in service 1 - subscriber sub-
21, sla-profile sla-profile-2 : host-limit ipv4-dhcp (1) exceeded "
```

Multi-sub-sap limit

The multi-sub-sap parameter defines the maximum number of subscribers (dynamic and static) that can be simultaneously active on this SAP. By default only a single subscriber is allowed (no multi-sub-sap).

Bridged CO @ BSA-1:

```
configure
    service
        vpls 1
            sap 1/1/3:2
                sub-sla-mgmt
                    multi-sub-sap 2
                exit
            exit
        exit
    exit
exit
```

Routed CO @ BSR-1:

```
configure
    service
        vprn 1
            subscriber-interface "sub-int-1"
                group-interface "group-int-1"
                    sap 1/1/3:2
                        sub-sla-mgmt
                            multi-sub-sap 2
                        exit
                    exit
                exit
            exit
        exit
    exit
exit
```

If the limit is reached, a new subscriber will be denied access, an event is generated (log-id 99) and the corresponding DHCP ack message is dropped:

```
101 2016/02/
24 09:35:10.95 CET WARNING: DHCP #2005 vprn1 Lease State Population Error
"Lease state table population error on SAP 1/1/3:2 in service 1 -
 Number of subscribers exceeds the configured multi-sub-sap limit (2)"
```

DHCP Host Connectivity Verification

Because the DHCP protocol does not have a keep-alive mechanism and IP address renewal is not frequent enough, alternative mechanisms are needed to track reachability of DHCP hosts.

The first alternative is called Subscriber Host Connectivity Verification (SHCV). A periodic unicast ARP is sent to the DHCP host. The connectivity test fails:

- If for X consecutive unicast ARP requests no ARP reply is received within the specified retry-timeout ([10 — 60] seconds, default 10). The number of retries (X-1) is specified by the retry-count ([2 — 29], default 2). Hence, at minimum 3 unicast ARP requests are send before connectivity is lost.
- If the ARP reply contains an inconsistent IP/MAC compared with the local DHCP lease state

  For a failed connectivity test, an event is raised and optionally the DHCP lease state is removed from the system: cleaning up of all related resources (e.g. anti-spoof table) and sending a DHCP release to the DHCP server. When ESM is enabled, the DHCP host also is removed.

The interval for the periodic checks can be configured between 1 and 6000 minutes. If not specified, the default value of 10 minutes will be used.

The maximum time for DHCP host connectivity loss detection in this case is:

```
( (host-connectivity-verify interval) + ( (retry-count) * (retry-timeout) ) )
```

Figure 163 **Subscriber Host Connectivity Verification**

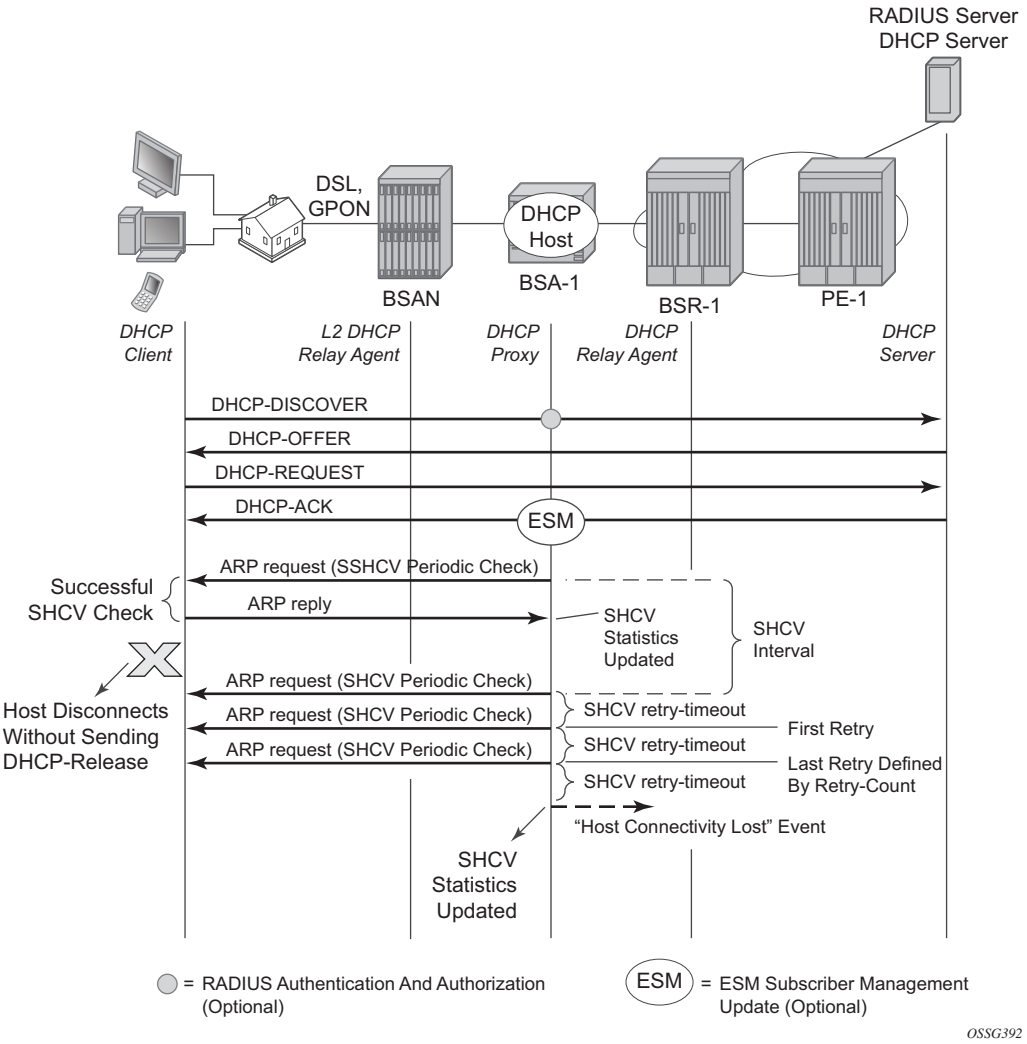


```
*A:BSA-1>config>service>vpls>sap# host-connectivity-verify ?
  - host-connectivity-verify source-ip <ip-address> [source-mac <ieee-
address>] [interval <interval>] [action {remove|alarm}] [timeout <retry-timeout>]
    [retry-count <count>]
 <ip-address>          : a.b.c.d
 <ieee-address>        : xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx
 <interval>            : [1..6000] minutes
 <{remove|alarm}>      : keywords
 <retry-timeout>       : [10..60] seconds
 <count>               : [2..29]
```

Bridged CO:

```
configure
    service
```

```
        vpls 1
            sap 1/1/3:2
                host-connectivity-verify source-ip 0.0.0.0 interval 1 action alarm
            exit
        exit
    exit
exit
```

The configured source IP should be an unused unique IP address in the DHCP client subnet or alternatively use source-ip 0.0.0.0. As the host-connectivity-verify application is sending a unicast ARP to the DHCP host, its ARP table is updated with the configured source-ip and source-mac (chassis MAC if not configured). If you would use an existing IP address, the DHCP host ARP table gets poisoned, breaking the connectivity to that host.

Routed CO:

```
configure
    service
        vprn 1
            subscriber-interface "sub-int-1"
                group-interface "group-int-1"
                    host-connectivity-verify interval 1 action remove
                exit
            exit
        exit
    exit
exit
```

The source IP is not configurable. The source-ip used in the unicast ARP is set to the local subscriber interface address in the subnet of the DHCP hosts that is checked for connectivity.

To verify the result of the connectivity check:

```
*A:BSA-1# show service id 1 host-connectivity-verify statistics
===============================================================================
Host connectivity check statistics
===============================================================================
SvcId   SapId/SdpId           HostIp
  DestIp                                            MAC
    Oper      Last-reply/Conn-lost
-------------------------------------------------------------------------------
1       1/1/3:2               10.1.0.101
  10.1.0.101                                        00:00:00:22:22:21
    Up        02/24/2016 09:38:29  (elapsed: 0d 00:00:22)
-------------------------------------------------------------------------------
1 host-connectivity states : 1 Up / 0 Down / 0 Retry pending
===============================================================================
*A:BSA-1#
```

With action alarm, the lease-state is not removed in case the connectivity with the host is lost. An event is generated (log-id 99) and the statistics show:

```
*A:BSA-1# show service id 1 host-connectivity-verify statistics
===============================================================================
Host connectivity check statistics
===============================================================================
SvcId   SapId/SdpId           HostIp
  DestIp                                            MAC
    Oper      Last-reply/Conn-lost
-------------------------------------------------------------------------------
1       1/1/3:2               10.1.0.101
  10.1.0.101                                        00:00:00:22:22:21
    Down      02/24/2016 15:56:36  (elapsed: 0d 00:02:38)
-------------------------------------------------------------------------------
1 host-connectivity states : 0 Up / 1 Down / 0 Retry pending
===============================================================================
*A:BSA-1#
```

In case the connectivity with the host is lost, following event is generated:

```
87 2016/02/24 15:58:15.98 CET WARNING: SVCMGR #2206 Base Host connectivity lost
"host connectivity lost on 1/1/
3:2 in service 1 for inetAddr = 10.1.0.101, chAddr=00:00:00:22:22:21."
```

When connectivity is restored, following event (log-id 99) is generated:

```
88 2016/02/24 15:59:35.98 CET WARNING: SVCMGR #2207 Base Host connectivity restored
"host connectivity restored on 1/1/
3:2 in service 1, for inetAddr = 10.1.0.101, chAddr=00:00:00:22:22:21."
```

Connectivity to a DHCP host can also be checked using an OAM command:

```
*A:BSR-1# oam host-connectivity-verify service 1 sap 1/1/3:2
===============================================================================
Triggering host connectivity verify for service 1 sap 1/1/3:2 ...
Waiting 3 seconds ...
===============================================================================
Host connectivity check statistics
===============================================================================
SvcId   SapId/SdpId           HostIp
  Oper    Last-reply/Conn-lost
    DestIp
-------------------------------------------------------------------------------
1       1/1/3:2               10.2.0.103
  Up      02/24/2016 16:35:10  (elapsed: 0d 00:00:02)
    10.2.0.103
-------------------------------------------------------------------------------
1 host-connectivity states : 1 Up / 0 Down / 0 Retry pending
===============================================================================
*A:BSR-1#
```

Note that in this case, no action is triggered. If the connectivity test is successful, the host-connectivity-verify statistics are updated with the new timestamp last-reply. If the connectivity test fails, the host-connectivity state becomes Retry Pending (oper state unknown) until an automatic test is scheduled again in the next interval.

To troubleshoot host-connectivity-verify, enable following debug log (additional filtering is possible on ip address, mac address and/or SAP):

```
debug
    service
        id 1
            host-connectivity-verify
            exit
        exit
    exit
exit
```

DHCP Lease Split

The second alternative is using a DHCP proxy server with the lease-split option.

A finer granularity of DHCP lease time is used between the DHCP client and the DHCP proxy server than between the DHCP proxy server and the DHCP server.

The maximum time for DHCP host connectivity loss detection in this case is the configured DHCP lease-split lease time.

DHCP communication is snooped between the DHCP client and DHCP server. In the DHCP ack message, the offered lease-time from the DHCP server is replaced with the configured DHCP proxy server lease-split lease time. Note that the lease time is only updated if the configured lease-split lease time is less than half of the original lease time value. The minimum value for the proxy server lease-split lease time is 5 minutes. When the DHCP client renews the DHCP session, the DHCP proxy server sends a DHCP ack on behalf of the DHCP server as long as the next renew time is earlier than half of the DHCP server expiry time for this session. With ESM enabled, RADIUS re-authentication will occur only when the DHCP request must be sent to the DHCP server. In other words, configuring a DHCP proxy with lease-split does not put extra load on the RADIUS server.

In the example in Figure 164, the DHCP server offers a lease time of 960 seconds. The lease time in the offer sent to DHCP client will be updated with the lease time of 300 seconds as configured in the DHCP proxy server lease-split on BSA-1.

Bridged CO @ BSA-1:

```
configure
    service
        vpls 1
            sap 1/1/3:2
                dhcp
                    proxy-server
                        lease-time min 5
                        no shutdown
```

```
                exit
            exit
        exit
    exit
exit
```

Routed CO @ BSR-1:

```
configure
    service
        vprn 1
            subscriber-interface "sub-int-1"
                group-interface "group-int-1"
                    dhcp
                        proxy-server
                            lease-time min 5
                            no shutdown
                        exit
                    exit
                exit
            exit
        exit
    exit
exit
```

Note: The emulated server address in the DHCP proxy-server configuration does not have to be configured for lease-split operation. This parameter is needed for an alternative use of the DHCP proxy server: RADIUS based IP configuration of a subscriber host. This is out of the scope of this configuration note.

If DHCP lease split is operational for a DHCP host, it will be shown in the Remaining Lifetime field of the detailed lease-state output. Note that the Session Timeout field is the original offered lease time from the DHCP server.

```
*A:BSA-1# show service id 1 dhcp lease-state detail
===============================================================================
DHCP lease states for service 1
===============================================================================
Service ID           : 1
IP Address           : 10.1.0.101
Client HW Address    : 00:00:00:22:22:21
SAP                  : 1/1/3:2
Termination Type     : local
Up Time              : 0d 00:00:26
Remaining Lease Time : 0d 00:04:34 (Lease Split)
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "sub-21"
Sub-Profile-String   : "sub-profile-2"
SLA-Profile-String   : "sla-profile-2"
App-Profile-String   : ""
Lease ANCP-String    : ""
```

```
Lease Int Dest Id    : ""
Category-Map-Name    : ""

Lease Info origin    : DHCP

Ip-Netmask           : 255.255.0.0
Broadcast-Ip-Addr    : 10.1.255.255
Default-Router       : 10.1.0.254
Primary-Dns          : N/A
Secondary-Dns        : N/A
Primary-Nbns         : N/A
Secondary-Nbns       : N/A

ServerLeaseStart     : 02/24/2016 09:51:00
ServerLastRenew      : 02/24/2016 09:51:00
ServerLeaseEnd       : 02/24/2016 21:51:00
Session-Timeout      : N/A
IPoE|PPP session     : No
Lease-Time           : 0d 12:00:00
DHCP Server Addr     : 172.16.0.1
Radius User-Name     : "00:00:00:22:22:21"
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
*A:BSA-1#
```

Figure 164 **DHCP Proxy Server: Lease Split Operation**

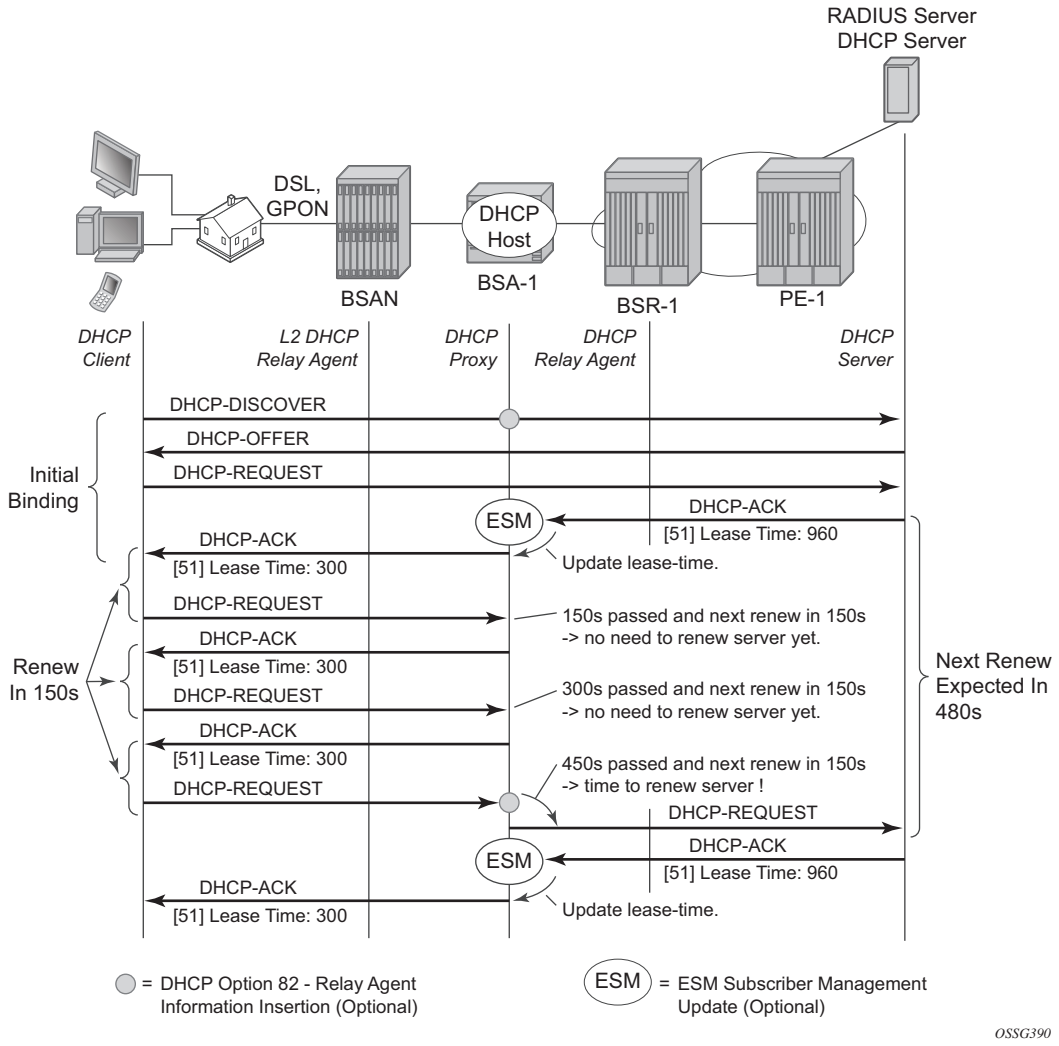

When the DHCP client disconnects without sending a DHCP release, the DHCP lease state in the BSA/BSR will be removed only when the DHCP lease time expires. With DHCP proxy server lease-split, the DHCP client disconnection can be sped up. In the example below, the DHCP client disconnection is detected in less than 5 minutes (lease-split lease time) while it would have taken up to 16 minutes without the lease-split. Note that the values are illustrative; in reality the DHCP lease times will be higher.

Figure 165 **DHCP Proxy Server: Lease Split Operation, DHCP Client Disconnected**

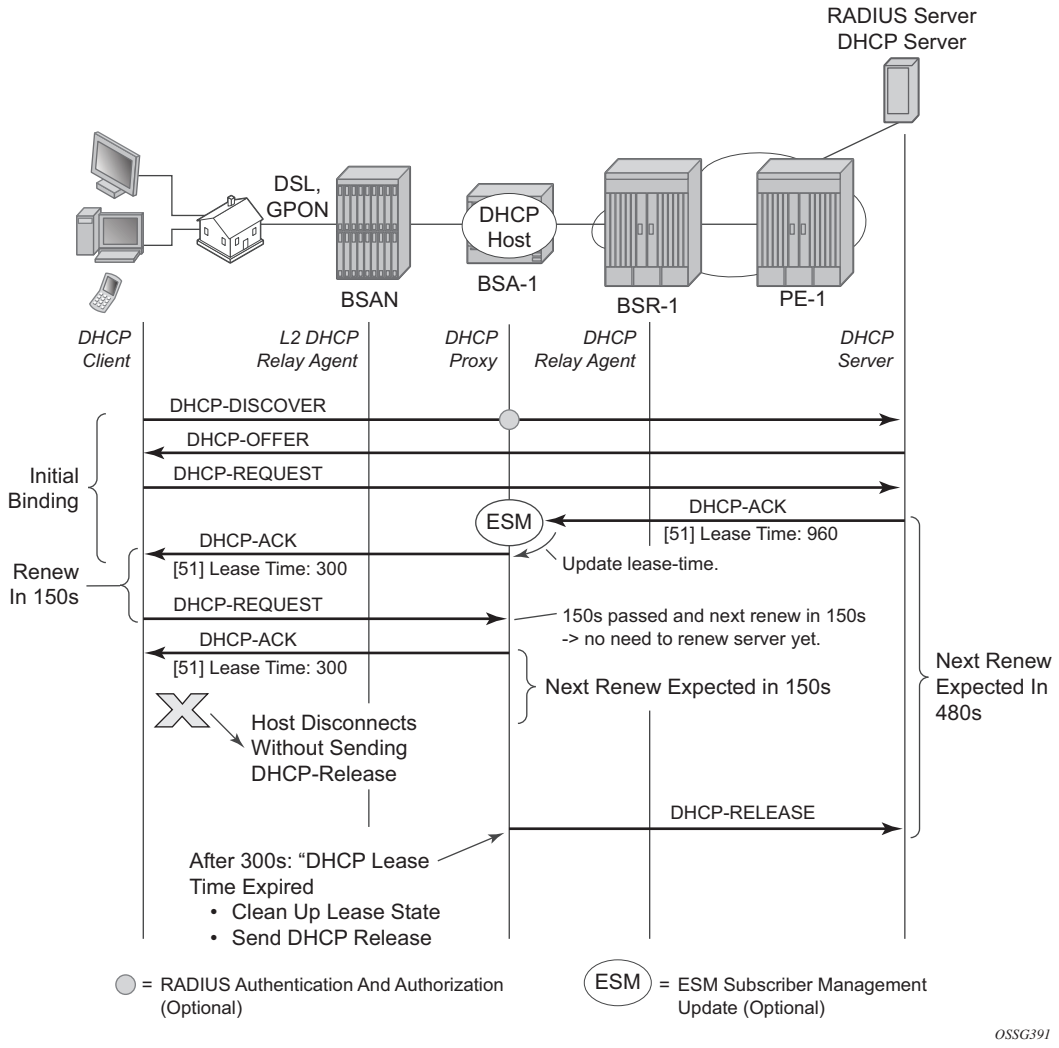

DHCP Host Mobility

A field technician verifying DSLAM operation often connects and disconnects from different ports rapidly. This will require the node to clear its own DHCP host state, the DHCP server state as well as flush MAC addresses learned within the VPLS network or clear ARP entries from the routing instance.

A DHCP request comes in on SAP2. On SAP1 there exists a lease state with the same Client Hardware address. The packet is dropped and a forced SHCV check verifies the existing lease state on SAP1. Three consecutive checks are launched with a timeout of 10 seconds. If the host indeed moved form SAP1 to SAP2, the connectivity check will fail on SAP1. The existing lease state is deleted and a DHCP release message is sent to the DHCP server. Any subsequent DHCP session setup will proceed as normal.

Figure 166 **DHCP Host Mobility**

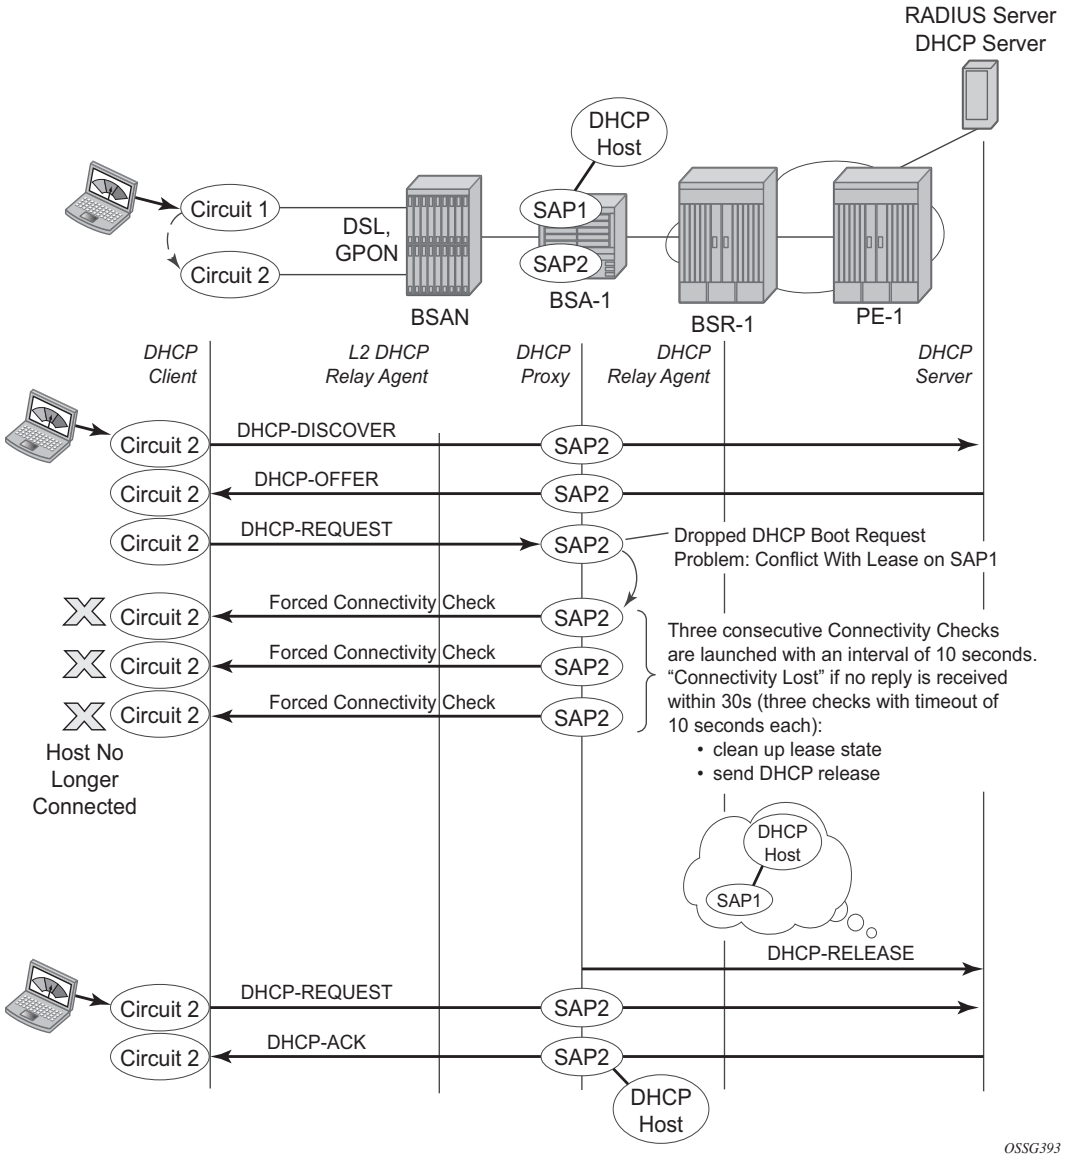

Note that for host mobility to function, host-connectivity-verification must be enabled. Next to periodic connectivity checks, it also enables forced checks triggered by moving hosts.

For Bridged CO, host-connectivity-verify must be enabled on the SAPs. When no interval is specified, it will default to 10 minutes for the periodic connectivity checks.

Bridged CO:

```
configure
    service
        vpls 1
            sap 1/1/3:1
                host-connectivity-verify source-ip 10.1.0.253
            exit
            sap 1/1/3:2
                host-connectivity-verify source-ip 10.1.0.253
            exit
        exit
    exit
exit
```

The configured source-ip should be an unused unique ip address in the DHCP client subnet or alternatively use source-ip 0.0.0.0. As the host-connectivity-verify application is sending a unicast ARP to the DHCP host, its ARP table is updated with the configured source-ip and source-mac (chassis MAC if not configured). If you would use an existing IP address, the DHCP host ARP table gets poisoned, breaking the connectivity to that host.

For Routed CO, host-connectivity-verify must be enabled on the group-interface. When no interval is specified, it will default to 10 minutes for the periodic connectivity checks.

Routed CO:

```
configure
    service
        vprn 1 customer 1 create
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    host-connectivity-verify
                exit
            exit
        exit
    exit
exit
```

The source IP address is not configurable. The source-ip used in the unicast ARP is fixed to the local subscriber interface address in the subnet of the DHCP hosts that is checked for connectivity.

Conclusion

This chapter provides configuration and troubleshooting commands for dynamic DHCP hosts. DHCP hosts can be instantiated in a Layer 2 bridged CO (VPLS) environment as well as in a Layer 3 Routed CO (IES/VPRN subscriber interface) context.

L2TP for Subscriber Access — LAC

This chapter provides information about L2TP for subscriber access.

Topics in this chapter include:

Applicability

This chapter describes L2TP Access Concentrator (LAC) support for the L2TP Aggregation Architecture (LAA) model and was initially written for SR OS release 11.0.R4. The CLI in the current edition is based on release 14.0.R2. PPP hosts are supported in a Routed CO model (with IES or VPRN services) using ATM, Ethernet or Ethernet over Pseudowire SAPs. A description of the L2TP Tunnel Switch (LTS) and L2TP Network Server (LNS) functions are out of the scope of this chapter.

Overview

PPP Access Architectures (PTA versus LAA)

The Broadband Forum proposes two architectures for Point-to-Point Protocol (PPP) access.

- The PPP Termination Aggregation Architecture (PTA)
- The L2TP Aggregation Architecture (LAA)

The PTA architecture (local-access model) uses the Broadband Network Gateway (BNG) to terminate user PPP sessions (see scenario PPPoE-5 in Figure 167).

The LAA architecture (which is a tunneled access model) uses a LAC to transport PPP sessions from the LAC to an LNS which performs tunnel termination (see scenario PPPoE-1 and PPPoE -2 in Figure 167).

Optionally, an LTS can be used in the transport network to perform the grooming of traffic between tunnels (see scenarios PPPoE-3 and PPPoE-4 in Figure 167).

The LNS is the logical termination point of the PPP sessions originated by the remote clients and tunneled by the LAC/LTS.

Figure 167 **Example Topology**

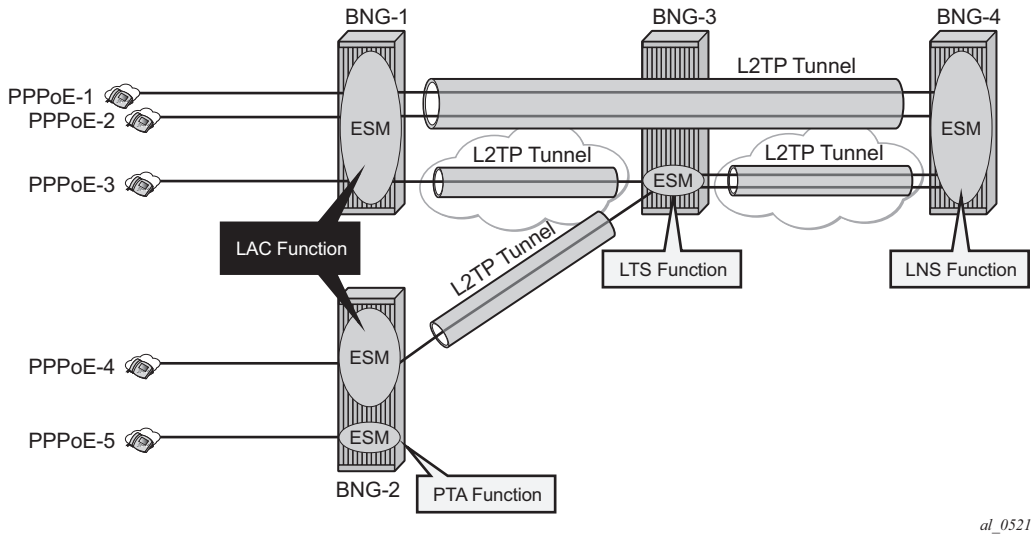

Supported L2TP Encapsulations

The router instance where the L2TP tunnel starts and where ESM is handled can be one and the same, but does not need to be the same. The LNS peer address can be reachable via IP, BGP/IGP shortcuts, over a spoke SDP (GRE/MPLS), RFC 4364 VPRNs (*BGP/MPLS IP Virtual Private Networks*), but cannot be an address belonging to a directly connected interface. See Figure 168.

Figure 168 **Supported LT2P Encapsulations**

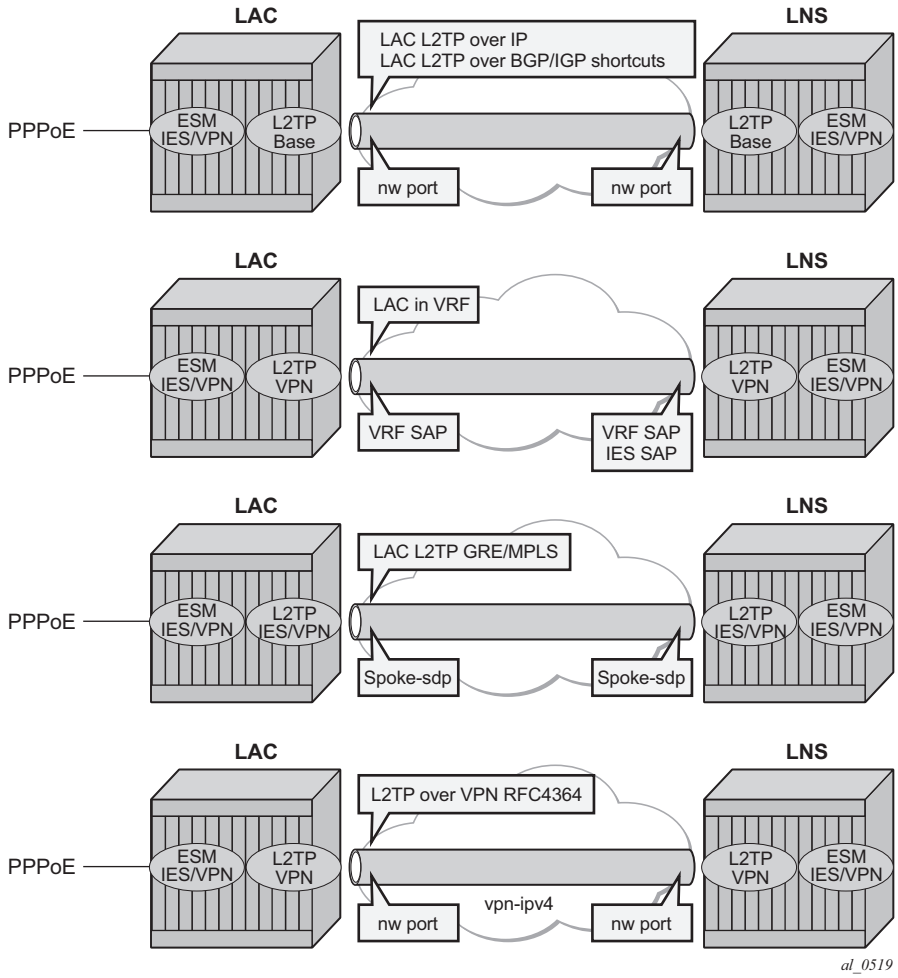

Recap of the L2TPv2 Protocol

L2TPv2 is a client-server protocol that encapsulates Layer 2 packets such as PPP, for transmission across a network and uses two different UDP message types:

- Control messages—L2TP passes control and data messages over separate control and data channels. The in-band control channel passes sequenced control connection management, call management, error reporting and session control messages. Optionally, a shared-secret challenge authentication method can be used between the tunnel endpoints. The following messages are used for L2TP tunnel session setup, teardown, and keepalive.
  - Tunnel setup (Control Connection Management)

- Start-Control-Connection-Request (SCCRQ)
- Start-Control-Connection-Reply (SCCRP)
- Start-Control-Connection-Connected (SCCCN)
- Stop-Control-Connection-Notification (StopCCN)

– Tunnel keepalive

- Hello (HELLO)

– Session setup (Call management) over an existing tunnel

- Incoming-Call-Request (ICRQ)
- Incoming-Call-Reply (ICRP)
- Incoming-Call-Connected (ICCN)
- Call-Disconnect-Notify (CDN)

The Zero-Length Body (ZLB) message is a control packet with only an L2TP header. ZLB messages are used for explicitly acknowledging packets on the reliable control channel.

To maximize extensibility while still permitting interoperability, a uniform method for encoding message types and bodies is used throughout L2TP via Attribute Value Pairs (AVP).

- Data messages — Data messages are used to encapsulate PPP frames that are sent into the L2TP tunnel.

L2TPv2 sessions run over an L2TP tunnel and are referenced by an L2TP session-id. An L2TP tunnel can carry none, one, or multiple L2TP sessions. An L2TP session corresponds to a PPPoE session. L2TPv3 for LAC-LNS dynamic tunnel setup is not supported.

L2TP Header and AVP Layout

The L2TPv2 header consists of following fields (RFC 2611, *URN Namespace Definition Mechanisms*):

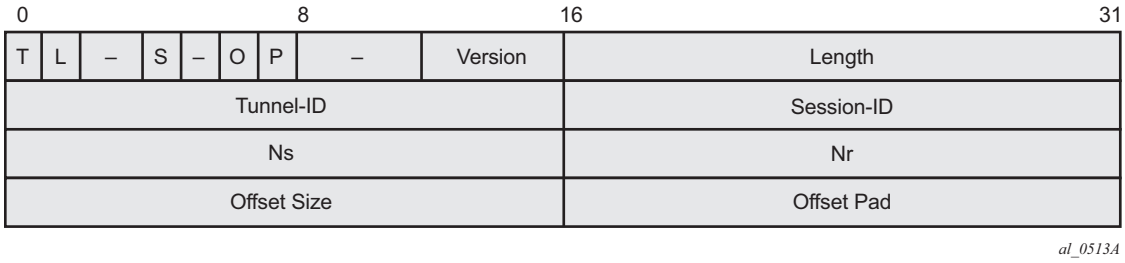

Table 35 **L2TPv2 Header Fields And Descriptions**

Field	Description
T	Type of L2TP message (1 bit): 0—data message 1—control message
L	Indicates if the optional Length field is present in the message (1 bit): 0—the field is left out of the message entirely 1—the field is included (must be included in control messages)
-	Reserved for future use, must be set to zero.
S	Indicates if the Ns and Nr fields are present (1 bit): 0 — the fields are left out of the message; entirely 1 — the fields are included (must be included in control messages)
O	Indicates if the Offset field is present (1 bit): 0 — the field is left out of the message entirely (must be left out of control messages); 1 — the field is included
P	Used with data messages only. Indicates priority of the data message (1 bit): 0 — no (this value is used for all control messages); 1 — yes
Version	The version of the message (4 bits): 2 — this is the latest version of the L2TP data message header; 1 — indicates an L2F packet as described in RFC 2341 Packets with an unknown version number are discarded.
Length	The total length (in bytes) of the L2TP message (16 bits).

Table 36 **L2TPv2 Fields And Descriptions**

Field	Description
Tunnel-ID	Identifies the L2TP tunnel (that is, the control connection).This number has local significance — each end gives the same tunnel different tunnel IDs. The ID refers to the receiver, not the sender, and is assigned during tunnel creation (16 bits).

Table 36 **L2TPv2 Fields And Descriptions (Continued)**

Field	Description (Continued)
Session-ID	Identifies the PPP session within a tunnel. This number has local significance — each end gives the same session different session IDs. The ID refers to the receiver, not the sender, and is assigned during session creation (16 bits).
Ns	The sequence number of the message. This is mandatory for control messages (to enable re-transmission of lost messages) but optional for data messages (to re-order data messages that were mis-sequenced during forwarding). The number, which starts at 0 and increments by 1, is assigned by an L2TP peer for each session in a tunnel (16 bits).
Nr	The sequence number of the next control message expected to be received. This is equal to the sequence number of last received control message plus 1. Used by the receiving peer to ensure that control messages are sent in order without duplication. In data messages, the field (if present as indicated by the S bit) is ignored (16 bits).
Offset Size	The location of the L2TP payload, expressed as the number of octets from the start of the message header (16 bits).
Offset pad	User-defined bytes used to pad the message header so that the payload starts at the location indicated by the Offset Size field (16 bits).

The AVP header consists of following fields (RFC2611):

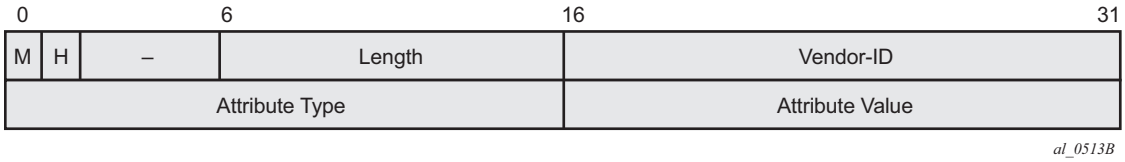

Table 37 **AVP Header Fields And Descriptions**

Field	Description
M	Mandatory bit — If the M bit is set on an unrecognized AVP within a message associated with a particular session, the session associated with this message MUST be terminated (1 bit).
H	Hidden bit — Identifies the hiding of data in the Attribute-Value field of an AVP. This capability can be used to avoid the passing of sensitive data, such as user passwords, as clear text in an AVP. The H-bit MUST only be set if a shared secret exists between the LAC and LNS. The shared secret is the same secret that is used for tunnel authentication. If the H-bit is set in any AVP(s) in a given control message, a Random Vector AVP must also be present in the message and MUST precede the first AVP having an H bit of 1 (1 bit).
-	Reserved for future use, must be set to zero (4 bits).
Length	Indicates the total number of bytes (including the overall length and bitmask fields) contained in this AVP (10 bits).

Table 37 **AVP Header Fields And Descriptions (Continued)**

Field	Description (Continued)
Vendor-id	Any vendor wishing to implement their own L2TP extensions can use their own Vendor ID along with private Attribute values. Vendor-ID=0 means that the standard AVPs are used (2 bytes).
Attribute Type	A value with a unique interpretation across all AVPs defined under a given Vendor (2 bytes).
Attribute Value	This is the actual value as indicated by the Vendor ID and Attribute Type (2 bytes).

RADIUS-Triggered Tunnel/Session Setup without LNS Renegotiation

Figure 169 depicts the complete PPP session setup, using RADIUS authentication on both LAC and LNS. After the discovery phase (PADI/PADO/PADR/PADS) and LCP negotiation phase (LCP config_request/Ack), the LAC initiates the L2TP tunnel setup based on Radius authentication information (Radius Request/Accept) and includes the negotiated PPP user-LAC information (called LCP proxy information). The LNS replies directly with a successful CHAP authentication if it agrees with the received proxy information. IP negotiation (IPCP config_request/Ack) is further handled between the user and the LNS and therefore the LAC has no IP knowledge of this PPP session.

Figure 169 **RADIUS Triggered Tunnel/Session Setup without LNS Renegotiation**

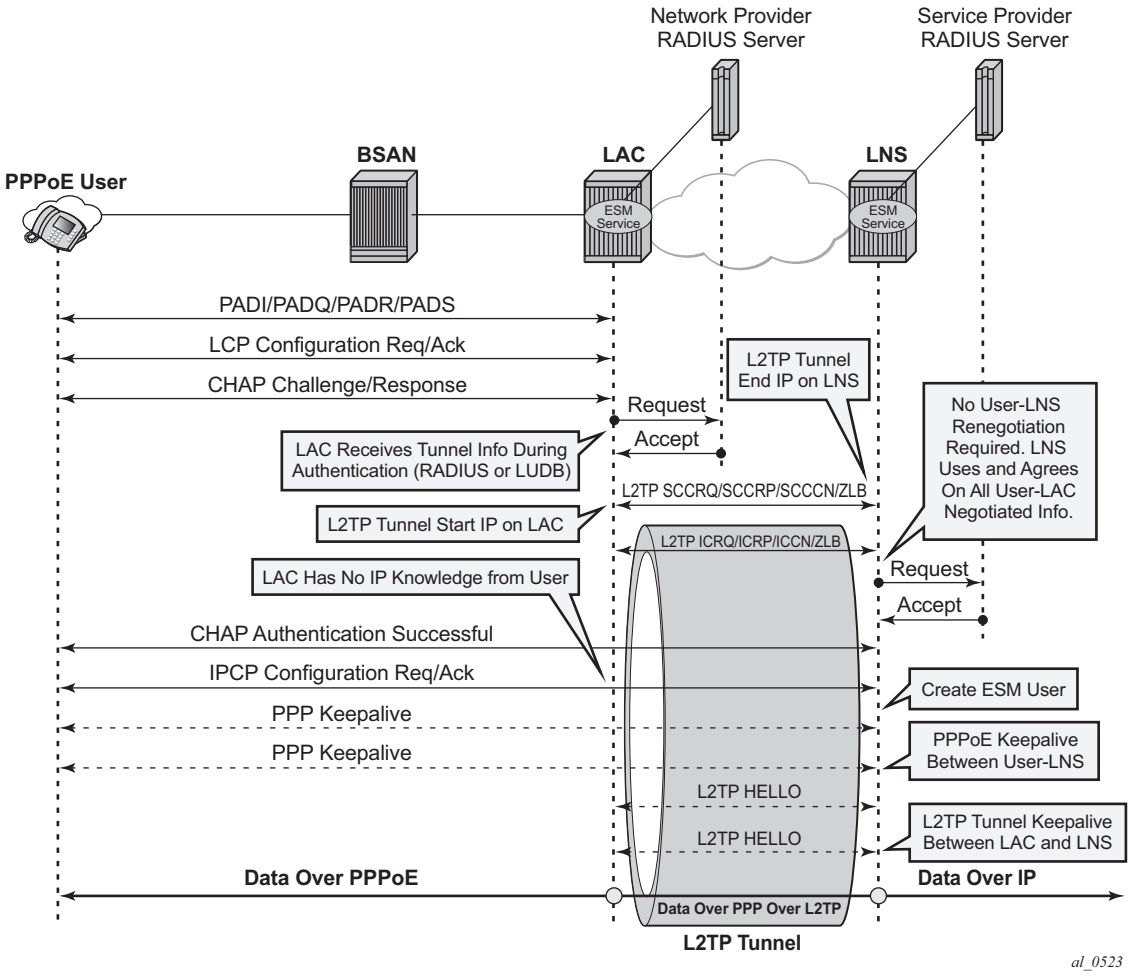

RADIUS-Triggered Tunnel/Session Setup with LNS Renegotiation

Figure 170 shows the scenario were the LNS does not agree with the received LCP proxy information and (re)starts the LCP phase (LCP config_request/Ack) directly with the PPP user. The rest of this scenario is the same as shown in Figure 169.

Figure 170 **RADIUS Triggered Tunnel/Session Setup with LNS Renegotiation**

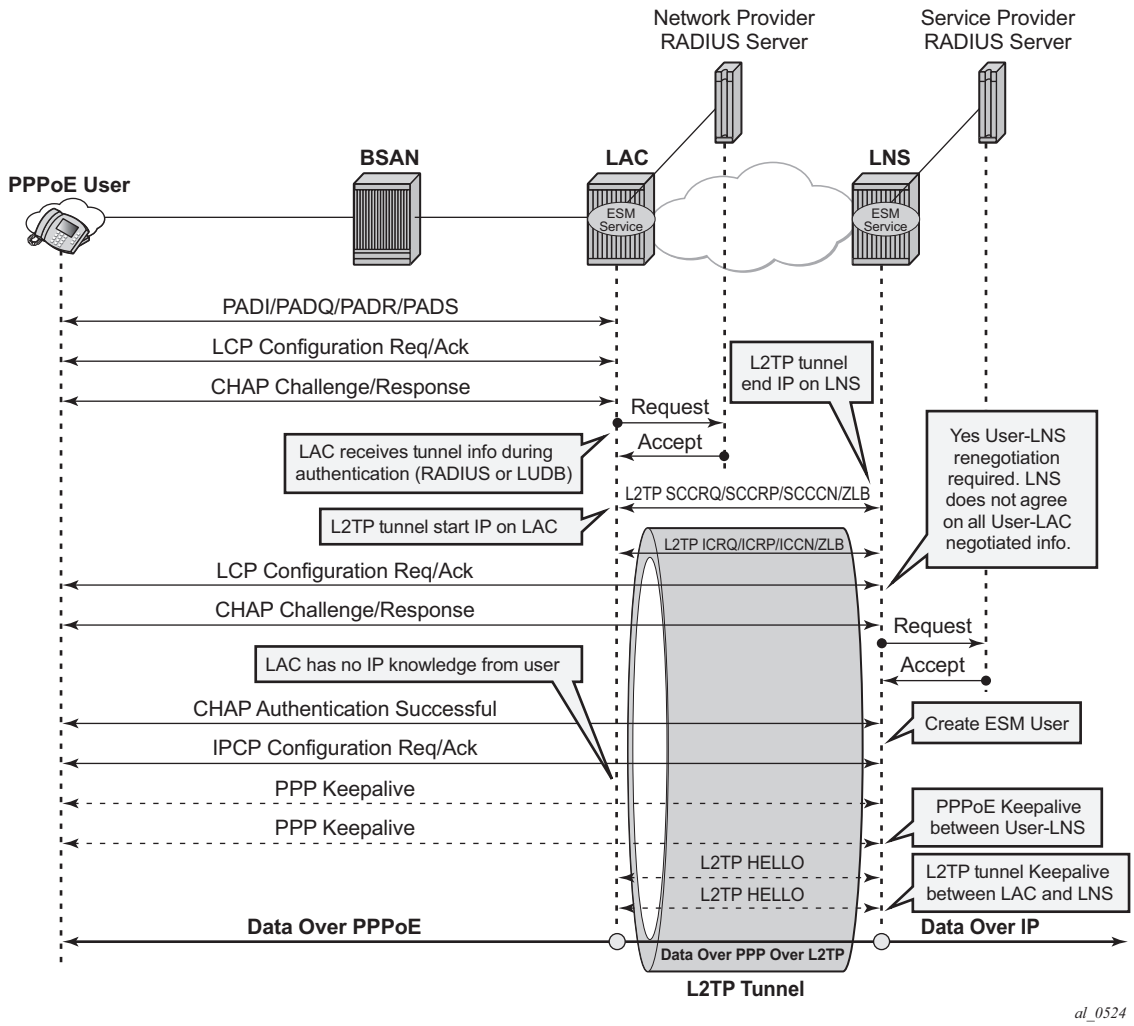

Running Multiple PPP Sessions Over a Single L2TP Tunnel

Figure 171 shows multiple PPP sessions tunneled over a single L2TP Tunnel. The LAC encapsulates each PPP session with a different L2TP session-id (SID) but with the same L2TP Tunnel-id (TID).

Figure 171 **Running Multiple PPP Sessions Over a Single L2TP Tunnel**

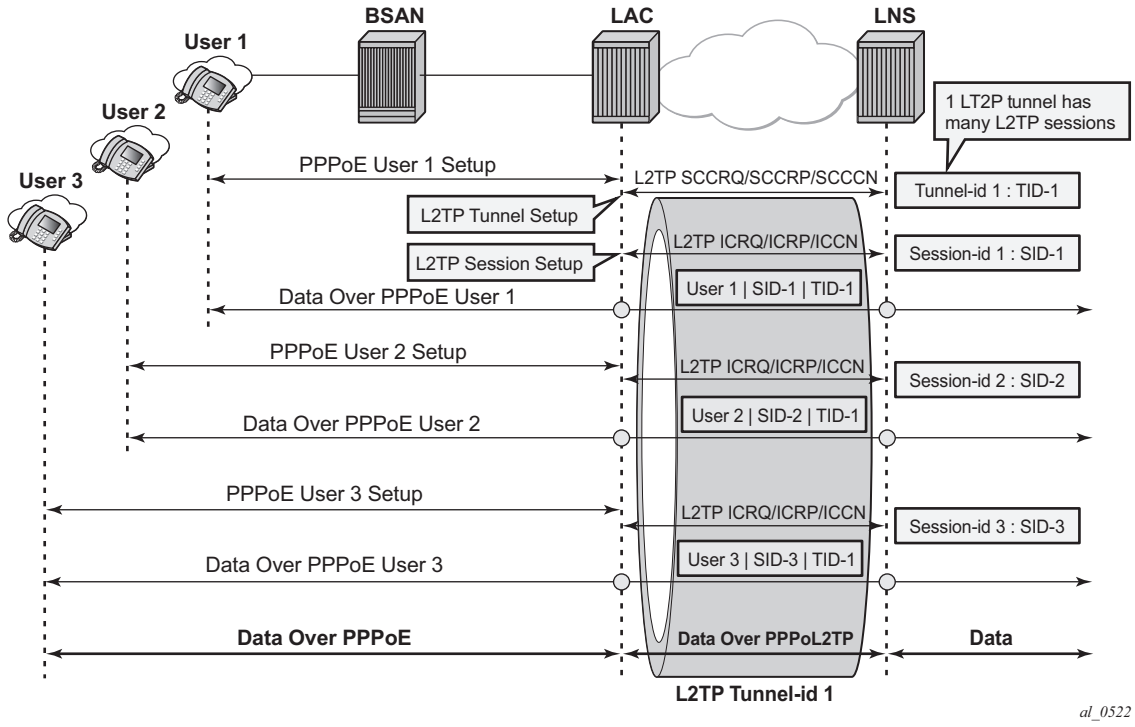

PPP User-Initiated Release/Terminate

Figure 172 shows the user initiated terminate_request tunneled by the LAC followed by the user initiated PADT terminated on the LAC. The LAC informs the LNS about the termination of the session via the L2TP CDN message. The L2TP tunnel can be optionally (idle-timeout) terminated via the L2TP StopCCN message.

Figure 172 **PPP User Initiated Release/Terminate**

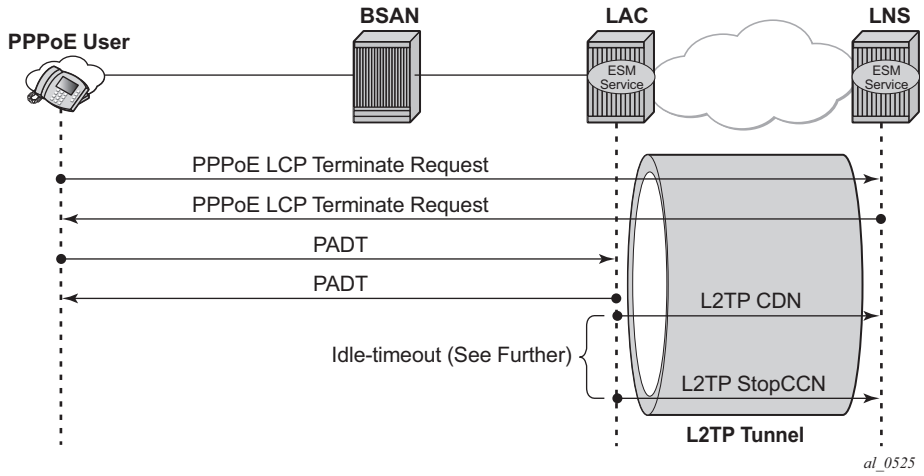

L2TP Tunnel/Session State Diagram

Figure 173 gives an overview of the main L2TP tunnel and session states. An L2TP tunnel in the establishedIdle state is a tunnel without sessions. A **tools** command (see Advanced Topics) can put an L2TP tunnel in a draining state (this prevents adding new sessions on tunnel but leaves the current sessions intact) or in a drained state (moved from draining to drained when all sessions terminated). The draining and drained state are not shown in the state diagram.

The L2TP tunnel setup occurs first with the triggers being: session activation, auto-establish, and a **tools start** command (see the advanced section). An L2TP session setup trigger is always session based.

Figure 173 **L2TP Tunnel and Session State Diagram**

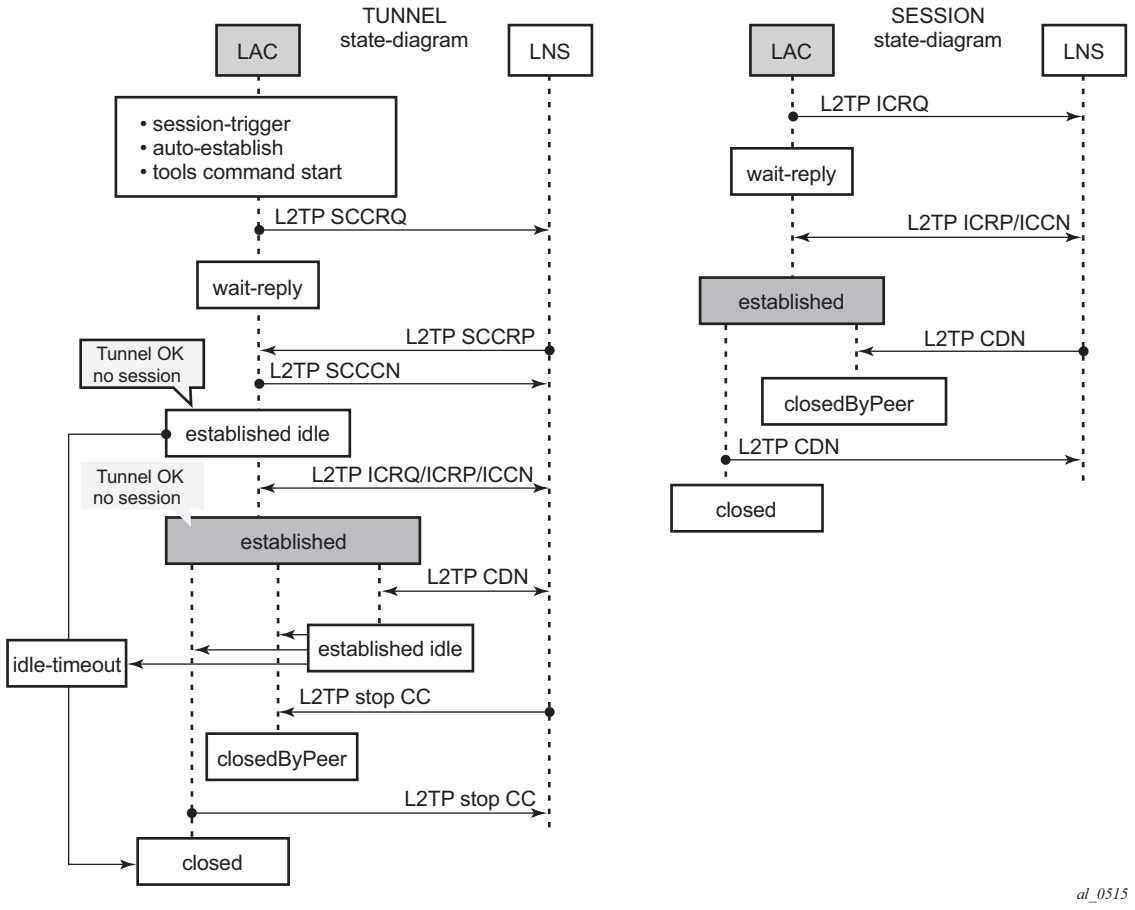

Configuration

Scenario 1: RADIUS-Derived L2TP Parameters

In the first scenario, the LAC receives an incoming connection and contacts the LAC RADIUS server. The RADIUS server retrieves the attributes for the user's domain (for example @wholesale.com) and passes the tunnel attributes to the LAC. Based on these RADIUS provided tunnel attributes, the LAC selects or initiates a new tunnel to the LTS or directly to the LNS. Once the tunnel is established, the LNS authenticates the end user using its own RADIUS server. Configuring the LNS and the LTS are out of the scope of this example.

In a RADIUS driven L2TP setup, either all or some of the required L2TP attributes are returned via RADIUS. If the RADIUS server returns only the L2TP [67] Tunnel-Server-Endpoint attributes, then the L2TP tunnel/session is established using the 'node parameter values' for the other required L2TP parameters. The 'l2tp node parameters' are defined under the configure router/service l2tp hierarchy. If the RADIUS server does not return all of the L2TP attributes and the node values are not configured, then the system falls back to default settings for these L2TP parameters.

The standard and vendor specific [26-6572] L2TP RADIUS attributes are listed in the tables below, together with the corresponding l2tp node parameters and defaults.

Table 38 **L2TP RADIUS Attributes**

Attribute ID	Attribute Name	Mandatory	CLI Node Parameter	Corresponding Defaults	
64	Tunnel-Type	Y	-	-	-
65	Tunnel-Medium-Type	Y	-	-	-
66	Tunnel-Client-Endpoint:[0-31]	N	local-address	no local-address	system-ip
67	Tunnel-Server-Endpoint	N	-	-	-
69	Tunnel-Password	N	password	no password	-
82	Tunnel-Assignment-ID:0	N	-	-	default_radius_group
82	Tunnel-Assignment-ID:[1..31]	N	-	-	Unnamed
83	Tunnel-Preference	N	preference	no preference	50
90	Tunnel-Client-Auth-ID	N	local-name	no local-name	system-name
91	Tunnel-Server-Auth-ID	N	-	-	-

Table 39 **L2TP RADIUS Attributes**

26-6527	Attribute Name	Mandatory	CLI Node Parameter	Corresponding Defaults	
-46	Alc-Tunnel-Group	N	-	-	-

Table 39 **L2TP RADIUS Attributes (Continued)**

26-6527	Attribute Name	Mandatory	CLI Node Parameter	Corresponding Defaults	
-47	Alc-Tunnel-Algorithm	N	session-assign-method	no session-assign-method	existingFirst
-48	Alc-Tunnel-Max-Sessions:0	N	-	group-session-limit	131071
-48	Alc-Tunnel-Max-Sessions:[1..31]	N	-	tunnel-session-limit	32767
-49	Alc-Tunnel-Idle-Timeout	N	idle-timeout	no idle-timeout	Infinite
-50	Alc-Tunnel-Hello-Interval	N	hello-interval	no hello-interval	300 sec
-51	Alc-Tunnel-Destruct-Timeout	N	destruct-timeout	no destruct-timeout	60 sec
-52	Alc-Tunnel-Max-Retries-Estab	N	max-retries-estab	no max-retries-estab	5
-53	Alc-Tunnel-Max-Retries-Not-Estab	N	max-retries-not-estab	no max-retries-not-estab	5
-54	Alc-Tunnel-AVP-Hiding	N	avp-hiding	no avp-hiding	Never
-97	Alc-Tunnel-Challenge	N	challenge	no challenge	Never
-104	Alc-Tunnel-Serv-Id	N	-	-	Base
-120	Alc-Tunnel-Rx-Window-Size	N	receive-window-size	no receive-window-size	64
-144	Alc-Tunnel-Acct-Policy	N	radius-accounting-policy	no radius-accounting-policy	-

LAC in the Base Routing Context (base) with Single Endpoint/ Single Tunnel

Using the mandatory L2TP RADIUS attributes (see the RADIUS user file below) the LAC establishes an L2TP tunnel. The source address for the tunnel is the IPv4address of a loopback interface in the Base router system (LAC tunnel endpoint).The destination for the tunnel is indicated by the Tunnel-Server-Endpoint RADIUS attribute [67], and is also known as the peer tunnel LNS endpoint address.

Figure 174 **LAC in Base Routing with Single Endpoint/Single Tunnel**

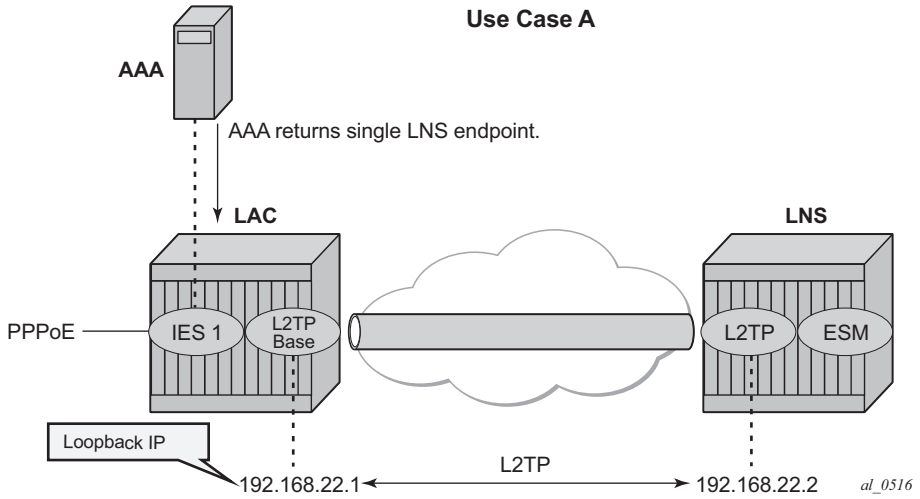

The PPPoE user terminates on IES service 1, sap 1/1/3:100, and is authenticated via RADIUS **authentication-policy** *authentication-1* which provides wholesale/retail (L2TP) information.

```
configure
    service
        ies 1 customer 1 create
            subscriber-interface "sub-l2tp" create
                unnumbered "system"
                group-interface "grp-l2tp" create
                    authentication-policy "radius-1"
                    sap 1/1/3:100 create
                        sub-sla-mgmt
                            sub-ident-policy "all-subscribers"
                            multi-sub-sap 1000
                            no shutdown
                        exit
                    exit
                    pppoe
                        sap-session-limit 10
                        no shutdown
                    exit
```

```
                exit
            exit
            no shutdown
        exit
    exit
exit
```

The excerpt from the FreeRADIUS users file below shows the attributes to be returned.

```
user1@wholesale.com     Cleartext-Password := "letmein", NAS-Identifier == "LAC"
                        Alc-Subsc-ID-Str = "%{User-name}",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Tunnel-Type:1 += L2TP,
                        Tunnel-Medium-Type:1 +=IP,
                        Tunnel-Server-Endpoint:1 += 192.168.22.2,
```

L2TP is enabled (no shutdown) in the related service instance.

The L2TP tunnel is set up in the base instance and not in a VRF because the attribute Alc-Tunnel-Serv-Id is not returned from RADIUS.

Missing L2TP parameters are taken from defaults defined in the router l2tp context.

```
configure router l2tp
       calling-number-format "%S %s"  # L2TP AVP 22 format
                                      # Default format 'system-name sap-id'
       ---snipped---
       no local-name                  # default name equals system-name
       no max-retries-estab           # default value equals 5
       ---snipped---
       no shutdown                    # enable L2TP
```

This scenario shows the PPPoE session termination (base IES service 1) and the L2TP tunnel setup in the base router instance.

LAC in the Base Routing Context with Multiple Endpoints

Figure 175 **LAC in the Base Routing with Multiple Endpoints**

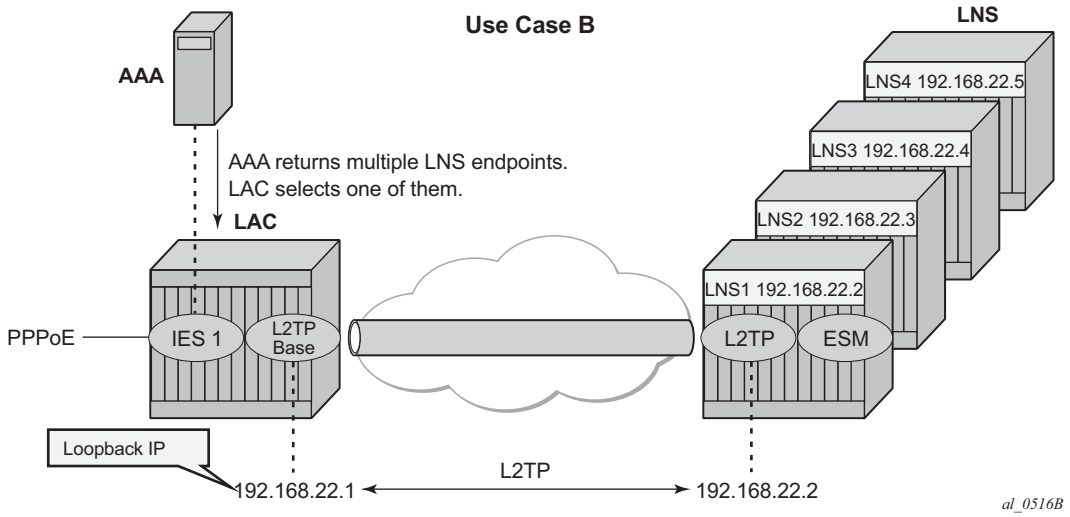

The following excerpt from the FreeRADIUS users file shows that user1@wholesale.com has 4 possible endpoints (LNS), each with its own tunnel preference. The LAC selects one L2TP endpoint out of these 4 tunnel specifications according to the configured L2TP selection process. This example uses weighted load balancing between LNS-T1 and LNS-T2 (tunnels LNS1-T1 and LNS2-T2 have best (lowest) and equal preference). The L2TP tunnel selection process is out of the scope of this example.

```
user1@wholesale.com     Cleartext-Password := "letmein", NAS-Identifier == "LAC"
                        Alc-Subsc-ID-Str = "%{User-name}",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
# group related info (tag 0)
                        Tunnel-Client-Endpoint:0 = 192.168.22.1,
                        Alc-Tunnel-Algorithm:0 = weighted-access,
                        Tunnel-Client-Auth-Id:0 = "lac-pe1",
                        Tunnel-Assignment-Id:0 = "L2TP-groupname",
                        Alc-Tunnel-Max-Retries-Estab:0 = 2,
# tunnel-1 related info (tag 1)
                        Tunnel-Type:1 += L2TP,
                        Tunnel-Medium-Type:1 +=IP,
                        Tunnel-Server-Endpoint:1 += 192.168.22.2,
                        Tunnel-Assignment-Id:1 += "LNS1-T1",
                        Tunnel-Preference:1 = 100,
# tunnel-2 related info (tag 2)
                        Tunnel-Type:2 += L2TP,
                        Tunnel-Medium-Type:2 +=IP,
                        Tunnel-Server-Endpoint:2 += 192.168.22.3,
                        Tunnel-Assignment-Id:2 += "LNS2-T2",
```

```
                          Tunnel-Preference:2 = 100,

# tunnel-3 related info (tag 3)
                        Tunnel-Type:3 += L2TP,
                        Tunnel-Medium-Type:3 +=IP,
                        Tunnel-Server-Endpoint:3 += 192.168.22.4,
                        Tunnel-Assignment-Id:3 += "LNS3-T3",
                        Tunnel-Preference:3 = 100,
--- snipped ---
```

This scenario also shows the PPPoE session termination (base IES service 1) and the L2TP tunnel setup in the base router instance.

LAC in a VRF

Figure 176 shows the PPPoE session termination (base IES service 1) and the L2TP tunnel setup in a different router instance (VPRN 65536).

Figure 176 **LAC in a VRF**

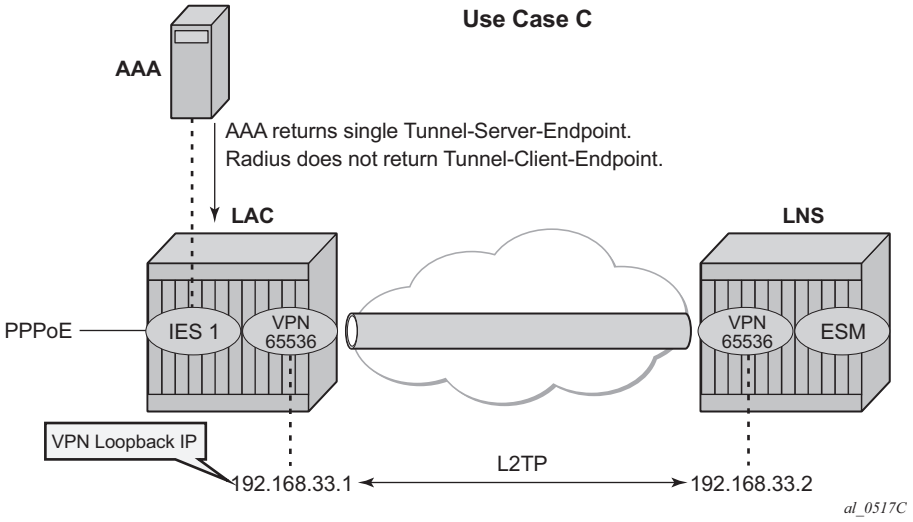

Using the following L2TP RADIUS attributes indicated, the LAC initiates an L2TP tunnel in VPRN 65536. The PPPoE session is still terminated in base IES service 1, which proves that both router instances can be different. (See use-case A for configuration details of IES service 1).

```
user1@wholesale.com     Cleartext-Password := "letmein", NAS-Identifier == "LAC"
                        Alc-Subsc-ID-Str = "%{User-name}",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-Tunnel-Serv-Id = 65536,
                        Tunnel-Client-Auth-Id:0 = "lac-pe1",
```

```
                        Tunnel-Assignment-Id:0 = "RADIUS-returned-TG",
                        Tunnel-Type:1 += L2TP,
                        Tunnel-Medium-Type:1 +=IP,
                        Tunnel-Server-Endpoint:1 += 192.168.33.2,
                        Tunnel-Assignment-Id:1 += "RADIUS-returned-TN",
```

If RADIUS does not return the L2TP source IP address (Tunnel-Client-Endpoint), then the IP address from the VPRN 65536 interface named ‘system’ is taken as the L2TP source address. The tunnel setup fails if this system interface is not created.

```
configure
    service
        vprn 65536
            interface "system" create
                address 192.168.33.1/32
                loopback
            exit
            l2tp
                no shutdown
            exit
        exit
    exit
exit
```

Scenario 2: Node-Derived L2TP Parameters

In the second scenario, the LAC receives the incoming connection and an ‘L2TP tunnel-group-name’ is assigned during LUDB or RADIUS authentication. This tunnel-group-name refers to the CLI preconfigured tunnel-group name context (**configure router <***router-name***> l2tp group <***tunnel-group-name***>**), which provides the context for all relevant tunnel attributes.

Based on these attributes, the LAC selects and initiates a tunnel to the LTS or directly to the LNS as in Scenario 1: RADIUS-Derived L2TP Parameters.

RADIUS Returns L2TP Group

In this use case (D), the L2TP tunnel-group-name is assigned during RADIUS authentication.

Figure 177 **RADIUS Returns L2TP Group**

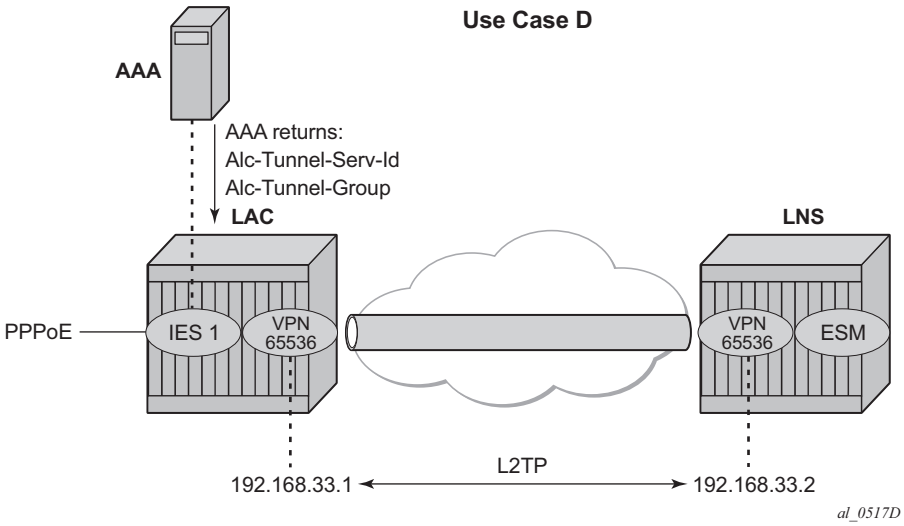


```
user1@wholesale.com     Cleartext-Password := "letmein", NAS-Identifier == "LAC"
                        Alc-Subsc-ID-Str = "%{User-name}",
                        Alc-Subsc-Prof-Str = "sub-profile-1",
                        Alc-SLA-Prof-Str = "sla-profile-1",
                        Alc-Tunnel-Serv-Id = 65536,
                        Alc-Tunnel-Group = "wholesale.com",
```

The L2TP tunnel is setup from VPRN 65536 (Alc-Tunnel-Serv-Id) and all L2TP tunnel information is taken from the l2tp group wholesale.com hierarchy (Alc-Tunnel-Group) as defined on the node.

```
configure
    service
        vprn 65536
            --- snipped ---
            interface "system" create
                address 192.168.33.1/32
                loopback
            exit
            l2tp
                group "wholesale.com" create
                    tunnel "wholesale.com" create
                        local-address 192.168.33.1
                        local-name "lac-pe1"
                        peer 192.168.33.2
                        no auto-establish
                        no shutdown
                    exit
                    no shutdown
                exit
                no shutdown
            exit
            no shutdown
        exit
```

```
    exit
exit
```

An L2TP tunnel is set up by either a PPP session-trigger, a **tools** command or by the l2tp group tunnel auto-establish parameter configuration. See the advanced section for the non-session-triggered tunnel setup.

RADIUS-Less Setup (LUDB Returns L2TP Group)

In this use case (E), the L2TP tunnel-group-name is assigned during LUDB authentication.

Figure 178 **RADIUS-Less Setup**

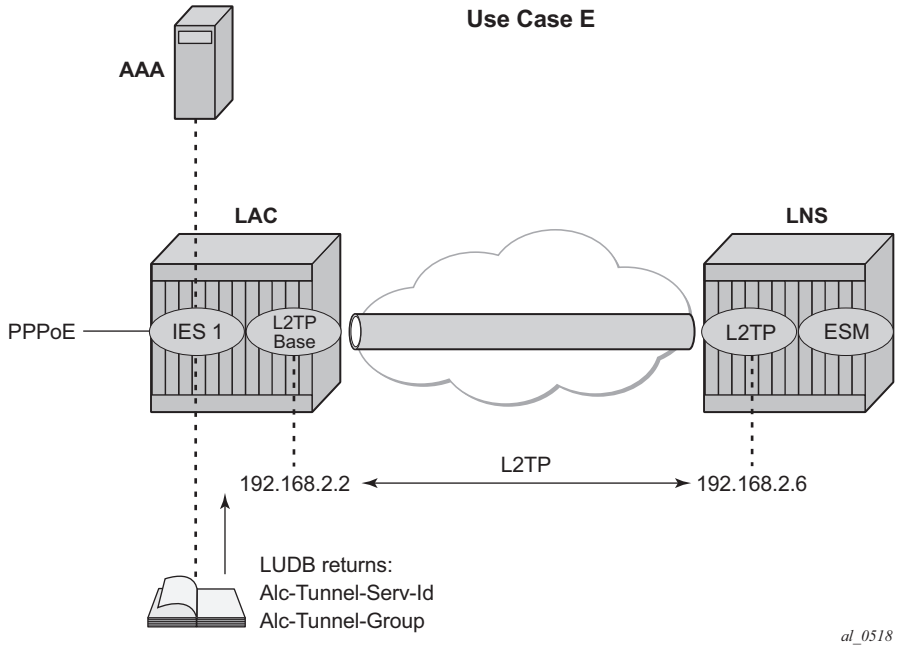

The PPPoE user enters on an IES service 1, sap 1/1/3:100, and is authenticated via the LUDB which provides L2TP wholesale/retail and ESM information. No RADIUS authentication is required in this example, because the PPPoE context refers to a local-user database *l2tp* to provide the subscriber authentication and the tunnel setup parameters.

```
configure
    service
        ies 1 customer 1 create
            subscriber-interface "sub-l2tp" create
                unnumbered "system"
```

```
                group-interface "grp-l2tp" create
                    sap 1/1/3:100 create
                        sub-sla-mgmt
                            sub-ident-policy "all-subscribers"
                            multi-sub-sap 1000
                            no shutdown
                        exit
                    exit
                    pppoe
                        sap-session-limit 10
                        user-db "l2tp"
                        no shutdown
                    exit
                exit
            exit
            no shutdown
        exit
    exit
exit
```

The referenced local user database *l2tp* configuration provides all of the required L2TP and ESM information.

```
configure
    subscriber-mgmt
        local-user-db "l2tp" create
            ppp
                match-list username
                host "wholesale.com" create
                    host-identification
                        username "wholesale.com" domain-only
                    exit
                    password ignore
                    identification-strings 254 create
                        subscriber-id "user@wholesale.com"
                        sla-profile-string "sla-profile-1"
                        sub-profile-string "sub-profile-1"
                    exit
                    l2tp
                        group "wholesale.com" service-id 65536
                    exit
                    no shutdown
                exit
            exit
            no shutdown
        exit
    exit
exit
```

Operation and Troubleshooting

This section explains how the use cases A to E described in the configuration section are verified using show, debug, and tools commands.

The standard router debugging tools can be used to monitor and troubleshoot the L2TP tunnel and session setup.

Overview of Debug and Show Commands

```
# useful show commands
show service id <service-id> ppp session [detail]
show router l2tp tunnel [detail]
show router l2tp session [detail]
show router l2tp peer [ip-address]

# debug to show PPPoE packet interaction
debug service id <service-id> ppp packet mode egr-ingr-and-dropped
debug service id <service-id> ppp packet detail-level medium

# debug to show RADIUS authentication interaction
debug router radius packet-type authentication

# debug to show LUDB authentication interaction
debug subscriber-mgmt local-user-db <local-user-db-name> detail all

# debug to show LAC Tunnel selection process and L2TP state-machine
debug router l2tp event lac-session-setup
debug router l2tp event finite-state-machine

# debug to show L2TP Tunnel and session setup
debug router l2tp packet direction both
debug router l2tp packet detail-level high
```

Understanding the Debug l2tp Command Output

The following L2TP ICRQ message extract (**debug router l2tp packet**) is used to explain how the displayed debug output should be interpreted. See Recap of the L2TPv2 Protocol -L2TP Header and AVP Layout for more details.

```
19 2016/06/03 13:44:32.96 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 8969 session 0, ns 2 nr 1, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallRequest(10)
    AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
        16795
    AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
        14268
    AVP CallingNumber(0,22), flags: mandatory, reserved=0
        "LAC 1/1/3:100""
```

- L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
  - version: v2

  - type field (T-bit): control message (ctrl)
  - 192.0.2.1:1701 -> 192.168.22.2:1701
    - 192.0.2.1:1701 - source tunnel-end-point:source udp port
    - 192.168.22.2:1701 - destination tunnel-end-point:destination udp port
- tunnel 8969 session 0, ns 2 nr 1, flags:, reserved=0
  - tunnel-id:8969
  - session-id:0
  - ns:2
  - nr:1
  - flags: 0 (refers to T/L/S/O/P bits L2TP header)
  - reserved field:0
- AVP CallingNumber(0,22), flags: mandatory, reserved=0
  - AVP MessageType(0,22): "LAC 1/1/3:100"
    - Vendor-id: 0 - Standard Attribute
    - Attribute Type: 22 – Calling Number AVP
    - Attribute Value: "LAC 1/1/3:100"

Scenario 1: RADIUS-Derived L2TP Parameters

LAC in Base Routing Context (base) with Single Endpoint/Single Tunnel

The ‘**debug service id <service-id> ppp packet mode egr-ingr-and-dropped** command shows PPPoE packet interaction. The following snapshot from the PADI packet shows the service, SAP, and received PPPoE tags. The received PPPoE DSL forum tags are by default copied during the LAC L2TP tunnel setup into the Incoming Call Request (ICRQ) DSL Forum AVP’s (RFC 5515).

```
1 2016/06/03 13:44:32.94 CEST MINOR: DEBUG #2001 Base PPPoE
"PPPoE: RX Packet
   IES 1, SAP 1/1/3:100

   DMAC: ff:ff:ff:ff:ff:ff
   SMAC: 00:00:00:00:01:01
   Ether Type: 0x8863 (Discovery)

   PPPoE Header:
   Version: 1                Type      : 1
   Code   : 0x09 (PADI)      Session-Id: 0x0000 (0)
   Length : 9

   PPPoE Tags:
```

```
   [0x0101] Service-Name: ""
   [0x0103] Host-Uniq: len = 1, value = 39
   [0x0105] Vendor-Specific: vendor-id = 0x0de9 (ADSL Forum)
      [0x01] Agent-Circuit-Id: "circuit0"
      [0x02] Agent-Remote-Id: "remote0"
      [0x81] Actual-Upstream: 2000
      [0x82] Actual-Downstream: 4000"
"
```

The **debug router radius packet-type authentication** command shows the actual authentication parameters returned by RADIUS. This example returns the minimum set of L2TP related RADIUS attributes.

```
12 2016/06/03 13:44:32.96 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 1 len 116 from 172.16.1.1:1812 vrid 1 pol rsp-radius-1
    VSA [26] 21 Alcatel(6527)
      SUBSC ID STR [11] 19 user1@wholesale.com
    VSA [26] 15 Alcatel(6527)
      SUBSC PROF STR [12] 13 sub-profile-1
    VSA [26] 15 Alcatel(6527)
      SLA PROF STR [13] 13 sla-profile-1
    TUNNEL TYPE [64] 4 1 L2TP(3)
    TUNNEL MEDIUM TYPE [65] 4 1 IPv4(1)
    TUNNEL SERVER ENDPOINT [67] 13 1 192.168.22.2
"
```

The **debug router l2tp event lac-session-setup** command shows the LAC tunnel selection for this example. An L2TP group-name '*default_radius_group*' with tunnel-name '*unnamed*' is created in this case.

```
13 2016/06/03 13:44:32.96 CEST MINOR: DEBUG #2001 Base PPPoE 1->L2TP
"PPPoE 1->L2TP: UDP 192.0.2.1:1701 -> 192.168.22.2:1701
preference 50 tunnel default_radius_group:unnamed
    request to open new tunnel 3897"

14 2016/06/03 13:44:32.96 CEST MINOR: DEBUG #2001 Base PPPoE 1->L2TP
"PPPoE 1->L2TP: UDP 192.0.2.1:1701 -> 192.168.22.2:1701
preference 50 tunnel default_radius_group:unnamed
    create session 255410587"
```

The '**debug router l2tp packet detail-level**' command shows the L2TP tunnel and session setup for this example.

Tunnel setup: The LAC sends a Start-Control-Connection-Request (SCCRQ) containing the assigned tunnel-id (no tunnel authentication in the example). The tunnel is now in a wait-reply state.

```
15 2016/06/03 13:44:32.96 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 0 session 0, ns 0 nr 0, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        StartControlConnectionRequest(1)
    AVP ProtocolVersion(0,2), flags: mandatory, reserved=0
```

```
        version=1, revision=0
    AVP HostName(0,7), flags: mandatory, reserved=0
        "lac-pe1"
    AVP WindowSize(0,10), flags: mandatory, reserved=0
        64
    AVP FramingCapabilities(0,3), flags: mandatory, reserved=0
        sync=no, async=no
    AVP BearerCapabilities(0,4), flags: mandatory, reserved=0
        digital=yes, analogue=no
    AVP FirmwareRevision(0,6), flags:, reserved=0
        3584
    AVP VendorName(0,8), flags:, reserved=0
        "Nokia"
    AVP AssignedTunnelId(0,9), flags: mandatory, reserved=0
        3897"
```

Tunnel setup: The LNS can bring up the tunnel, so the LNS replies with a Start-Control-Connection-Reply (SCCRP) including the assigned tunnel-id.

```
16 2016/06/03 13:44:32.96 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.22.2:1701 -> 192.0.2.1:1701
tunnel 3897 session 0, ns 0 nr 1, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        StartControlConnectionReply(2)
    AVP ProtocolVersion(0,2), flags: mandatory, reserved=0
        version=1, revision=0
    AVP HostName(0,7), flags: mandatory, reserved=0
        "lns-pe2"
    AVP WindowSize(0,10), flags: mandatory, reserved=0
        64
    AVP FramingCapabilities(0,3), flags: mandatory, reserved=0
        sync=no, async=no
    AVP BearerCapabilities(0,4), flags: mandatory, reserved=0
        digital=yes, analogue=no
    AVP FirmwareRevision(0,6), flags:, reserved=0
        3584
    AVP VendorName(0,8), flags:, reserved=0
        "Nokia"
    AVP AssignedTunnelId(0,9), flags: mandatory, reserved=0
        8969"
```

Tunnel setup: The LAC responds with a Start-Control-Connection-Connected (SCCCN) message. After an LNS ZLB acknowledgment, the tunnel is in the establishedIdle state.

```
17 2016/06/03 13:44:32.96 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 8969 session 0, ns 1 nr 1, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        StartControlConnectionConnected(3)"
```

Session setup: Once the tunnel exists, a three-way exchange for session establishment within the tunnel is performed. The LAC sends an Incoming-Call-Request (ICRQ) with the parameter information for the session. The session is now in the wait-reply state.

```
19 2016/06/03 13:44:32.96 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 8969 session 0, ns 2 nr 1, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallRequest(10)
    AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
        16795
    AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
        14268
    AVP CallingNumber(0,22), flags: mandatory, reserved=0
        "LAC 1/1/3:100"
    AVP AgentCircuitId(3561,1), flags:, reserved=0
        "circuit0"
    AVP AgentRemoteId(3561,2), flags:, reserved=0
        "remote0"
    AVP ActDataRateUp(3561,129), flags:, reserved=0
        2000000
    AVP ActDataRateDown(3561,130), flags:, reserved=0
        4000000""
```

Session setup: The LNS sends an Incoming-Call-Reply (ICRP) that contains the assigned session-id. The session is now in the connect state.

```
21 2016/06/03 13:44:32.96 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, ingress)
"L2TP(v2, ctrl, ingress): UDP 192.168.22.2:1701 -> 192.0.2.1:1701
tunnel 3897 session 16795, ns 1 nr 3, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallReply(11)
    AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
        5378"
```

Session setup: The LAC sends an Incoming Call Connected (ICCN) and provides the LNS with additional information from the user initiated session. This information includes the LCP information from the negotiation that the LAC and remote user performed. This information is used by the LNS to decide whether to start LCP re-negotiation and/or Authentication re-negotiation with the PPP user or not. After an LNS ZLB acknowledgment the session is in the established state.

```
24 2016/06/03 13:44:32.96 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 8969 session 5378, ns 3 nr 2, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallConnected(12)
    AVP FramingType(0,19), flags: mandatory, reserved=0
        sync=no, async=no
    AVP TxConnectSpeed(0,24), flags: mandatory, reserved=0
        4000000
    AVP InitialRxLcpConfReq(0,26), flags:, reserved=0
         01 04 05 d4
         [1] MRU: 1492
    AVP LastTxLcpConfReq(0,27), flags:, reserved=0
         01 04 05 d4 03 05 c2 23 05 05 06 6f 9e d4 4b
         [1] MRU: 1492
         [3] Authentication-Protocol: 0xc223 (CHAP), Algorithm = 5 (MD5)
         [5] Magic-Number: 0x6f9ed44b
    AVP LastRxLcpConfReq(0,28), flags:, reserved=0
```

```
         01 04 05 d4
         [1] MRU: 1492
    AVP ProxyAuthenType(0,29), flags:, reserved=0
        chap(2)
    AVP ProxyAuthenName(0,30), flags:, reserved=0
         "user1@wholesale.com"
    AVP ProxyAuthenChallenge(0,31), flags:, reserved=0
         b8 ed 5f f5 c2 3e 18 ec bd 8f 48 83 db 56 51 cb
         f6 f5 4b 06 9f 09 34 2b ac ea 6f da 49 1c df 03
         8a 40 78 cc 7e 10 ba bc 9f 82 40 fa d8 11 47 cf
         07 92 d6 26 1b 8a d0 47 f4
    AVP ProxyAuthenId(0,32), flags:, reserved=0
        id=1, reserved=0
    AVP ProxyAuthenResponse(0,33), flags:, reserved=0
         db e1 a4 16 c6 5e 33 11 35 57 3b 25 04 aa 71 18
    AVP RxConnectSpeed(0,38), flags:, reserved=0
        2000000"
```

The PPPoE session operational information for IES 1/base instance is as follows.

```
*A:LAC# show service id 1 ppp session
===============================================================================
PPP sessions for service 1
===============================================================================
User-Name
  Descr.
           Up Time       Type  Termination     IP/L2TP-Id/Interface-Id MC-Stdby
-------------------------------------------------------------------------------
user1@wholesale.com
  svc:1 sap:1/1/3:100 mac:00:00:00:00:01:01 sid:1
           0d 00:00:32   oE    lac             255410587
-------------------------------------------------------------------------------
No. of PPP sessions: 1
===============================================================================
*A:LAC#
```

The tunnel operational information in base instance shows that the tunnel is established.

```
*A:LAC# show router l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State              Blacklist-state   Ses Active
  Group                                                             Ses Total
    Assignment
-------------------------------------------------------------------------------
255393792  3897      8969      established        not-blacklisted   1
  default_radius_group                                              1

    unnamed
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
*A:LAC#
```

Detailed tunnel operational information is obtained using following command.

```
*A:LAC# show router l2tp tunnel detail
===============================================================================
L2TP Tunnel Status
===============================================================================

Connection ID: 255393792
State        : established
IP           : 192.0.2.1
UDP          : 1701
Peer IP      : 192.168.22.2
Peer UDP     : 1701
Tx dst-IP    : 192.168.22.2
Tx dst-UDP   : 1701
Rx src-IP    : 192.168.22.2
Rx src-UDP   : 1701
Name         : lac-pe1
Remote Name  : lns-pe2
Assignment ID: unnamed
Group Name   : default_radius_group
Acct. Policy : N/A
Error Message: N/A

                                       Remote Conn ID    : 587792384
Tunnel ID         : 3897               Remote Tunnel ID  : 8969
Preference        : 50                 Receive Window    : 64
Hello Interval (s): 300
Idle TO (s)       : infinite           Destruct TO (s)   : 60
Max Retr Estab    : 5                  Max Retr Not Estab: 5
Session Limit     : 32767              AVP Hiding        : never
Transport Type    : udpIp              Challenge         : never
Time Started      : 06/03/2016 13:44:33 Time Idle        : N/A
Time Established  : 06/03/2016 13:44:33 Time Closed      : N/A
Stop CCN Result   : noError            General Error     : noError
Blacklist-state   : not-blacklisted
Set Dont Fragment : true

Failover
State             : not-recoverable
Recovery Conn ID  : N/A
Recovery state    : not-applicable
Recovered Conn ID : N/A
Recovery method   : mcs
Track SRRP        : (Not specified)
Ctrl msg behavior : handle
Recovery time (ms)
Requested         : N/A
Peer              : N/A
-------------------------------------------------------------------------------

No. of tunnels: 1
===============================================================================
*A:LAC#
```

The session operational information shows the session is established.

```
*A:LAC# show router l2tp session

===============================================================================
```

```
L2TP Session Summary
===============================================================================
ID                  Control Conn ID     Tunnel-ID   Session-ID  State
-------------------------------------------------------------------------------
255410587           255393792           3897        16795       established
-------------------------------------------------------------------------------
No. of sessions: 1
===============================================================================
*A:LAC#
```

For detailed operational session information use the following command.

```
*A:LAC# show router l2tp session detail

===============================================================================
L2TP Session 255410587
===============================================================================

Connection ID: 255410587
State        : established
Tunnel Group : default_radius_group
Assignment ID: unnamed
Error Message: N/A

Control Conn ID   : 255393792           Remote Conn ID    : 587797762
Tunnel ID         : 3897                Remote Tunnel ID  : 8969
Session ID        : 16795               Remote Session ID : 5378
Time Started      : 06/03/2016 13:44:33
Time Established  : 06/03/2016 13:44:33 Time Closed       : N/A
CDN Result        : noError             General Error     : noError
-------------------------------------------------------------------------------

No. of sessions: 1
===============================================================================
*A:LAC#
```

LAC in Base Routing Context with Multiple Endpoints (Tunnel Selection Process)

The **debug router radius packet-type authentication** command shows the actual RADIUS authentication parameters returned. This example returns multiple tunnel endpoints from which the LAC selects one. This example uses weighted load balancing. (The L2TP tunnel selection process is out of the scope of this example).

```
40 2016/06/03 13:47:36.66 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 2 len 232 from 172.16.1.1:1812 vrid 1 pol rsp-radius-1
    VSA [26] 21 Alcatel(6527)
      SUBSC ID STR [11] 19 user1@wholesale.com
    VSA [26] 15 Alcatel(6527)
      SUBSC PROF STR [12] 13 sub-profile-1
    VSA [26] 15 Alcatel(6527)
      SLA PROF STR [13] 13 sla-profile-1
    TUNNEL CLIENT ENDPOINT [66] 12 192.168.22.1
```

```
    VSA [26] 6 Alcatel(6527)
      TUNNEL ALGORITHM [47] 4 weighted access(1)
    TUNNEL CLIENT AUTH ID [90] 7 lac-pe1
    TUNNEL ASSIGNMENT ID [82] 14 L2TP-groupname
    VSA [26] 6 Alcatel(6527)
      TUNNEL MAX RETRIES ESTAB [52] 4 0 2
    TUNNEL TYPE [64] 4 1 L2TP(3)
    TUNNEL MEDIUM TYPE [65] 4 1 IPv4(1)
    TUNNEL SERVER ENDPOINT [67] 13 1 192.168.22.2
    TUNNEL ASSIGNMENT ID [82] 8 1 LNS1-T1
    TUNNEL PREFERENCE [83] 4 1 100
    TUNNEL TYPE [64] 4 1 L2TP(3)
    TUNNEL MEDIUM TYPE [65] 4 1 IPv4(1)
    TUNNEL SERVER ENDPOINT [67] 13 1 192.168.22.3
    TUNNEL ASSIGNMENT ID [82] 8 1 LNS2-T2
"
```

The **debug router l2tp event lac-session-setup** command shows the LAC tunnel LNS2-T2 is selected for this example.

```
41 2016/06/03 13:47:36.65 CEST MINOR: DEBUG #2001 Base PPPoE 2->L2TP
"PPPoE 2->L2TP: UDP 192.168.22.1:1701 -> 192.168.22.2:1701
preference 100 tunnel L2TP-groupname:LNS2-T2
    request to open new tunnel 3896"

42 2016/06/03 13:47:36.65 CEST MINOR: DEBUG #2001 Base PPPoE 2->L2TP
"PPPoE 2->L2TP: UDP 192.168.22.1:1701 -> 192.168.22.2:1701
preference 100 tunnel L2TP-groupname:LNS2-T2
    create session 255353413"
```

The operational PPPoE session information in IES 1/base instance is shown as follows.

```
*A:LAC# show service id 1 ppp session

===============================================================================
PPP sessions for service 1
===============================================================================
User-Name
  Descr.
          Up Time       Type  Termination     IP/L2TP-Id/Interface-Id MC-Stdby
-------------------------------------------------------------------------------
user1@wholesale.com
  svc:1 sap:1/1/3:100 mac:00:00:00:00:01:01 sid:1
          0d 00:00:58   oE    lac             255353413
-------------------------------------------------------------------------------
No. of PPP sessions: 1
===============================================================================
*A:LAC#
```

The operational tunnel information (base instance) is shown below.

```
*A:LAC# show router l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State              Blacklist-state   Ses Active
  Group                                                             Ses Total
```

```
    Assignment
-------------------------------------------------------------------------------
255328256  3896      5500      established        not-blacklisted   1
  L2TP-
groupname                                                  1
    LNS2-T2
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
*A:LAC#
```

Operational session information (base instance) shows the session is in the established state.

```
*A:LAC# show router l2tp session

===============================================================================
L2TP Session Summary
===============================================================================
ID                  Control Conn ID     Tunnel-ID   Session-ID  State
-------------------------------------------------------------------------------
255353413           255328256           3896        25157       established
-------------------------------------------------------------------------------
No. of sessions: 1
===============================================================================
*A:LAC#
```

The L2TP endpoint/peer information shows there are two tunnels for tunnel endpoint 192.168.22.2.

```
*A:LAC# show router l2tp peer
===============================================================================
L2TP Peers
===============================================================================
Peer IP                                          Port  Tun Active Ses Active
                                      Drain Reachability Tun Total  Ses Total
-------------------------------------------------------------------------------
192.168.22.2                                     1701  2          1
                                                       2          1
192.168.22.3                                     1701  0          0
                                                       0          0
-------------------------------------------------------------------------------
No. of peers: 2
===============================================================================
*A:LAC#
```

The following command gives a system overview of subscriber session related data. This system overview shows the current and peak values per session type (local PTA, LAC, LTS, LNS) and an overview of the number of originated or terminated L2TP tunnels. Peak values can be cleared via the **clear subscriber-mgmt peakvalue-stats** command.

```
*A:LAC# show subscriber-mgmt statistics session system

===============================================================================
```

```
Subscriber Management Statistics for System
===============================================================================
       Type                                Current     Peak      Peak Timestamp
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
PPP Session Statistics
-------------------------------------------------------------------------------
Local  PPP Sessions     - PPPoE                 0        0
       PPP Sessions     - PPPoEoA               0        0
       PPP Sessions     - PPPoA                 0        0
       PPP Sessions     - L2TP (LNS)            0        0
-------------------------------------------------------------------------------
LAC    PPP Sessions     - PPPoE                 1        1 06/03/2016 13:47:37
       PPP Sessions     - PPPoEoA               0        0
       PPP Sessions     - PPPoA                 0        0
       PPP Sessions     - L2TP (LTS)            0        0
-------------------------------------------------------------------------------
Total  PPP Sessions     - established           1        1 06/03/2016 13:47:37
       PPP Sessions     - in setup              0        1 06/03/2016 13:47:37
       PPP Sessions     - local                 0        0
       PPP Sessions     - LAC                   1        1 06/03/2016 13:47:37
-------------------------------------------------------------------------------
L2TP   L2TP Tunnels     - originator            2        2 06/03/2016 13:47:37
       L2TP Tunnels     - receiver              0        0
       Total L2TP Tunnels                       2        2 06/03/2016 13:47:37
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
IPOE Session Statistics
-------------------------------------------------------------------------------
Total  IPOE Sessions    - established           0        0
       IPOE Sessions    - in setup              0        0
-------------------------------------------------------------------------------
===============================================================================
Peak values last reset at : n/a
*A:LAC#
```

LAC in a VRF

This example returns VPRN 65536 as the L2TP service instance [26-6527-104 Alc-Tunnel-Serv-Id]. The VPRN 65536 interface system address is used as the L2TP source address since the attribute Tunnel-Client-Endpoint is not returned.

The ip-address 192.168.33.1 (Tunnel-Server-Endpoint) needs to be routable in VRF 65536 over a SAP or to a remote PE. This example uses BGP/MPLS IP Virtual Private Networks (VPNs) (RFC4364) to access the remote PE.

```
*A:LAC# show router 65536 route-table

===============================================================================
Route Table (Service: 65536)
===============================================================================
Dest Prefix[Flags]                            Type    Proto     Age        Pref
```

```
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
--- snipped ---
192.168.33.1/32                              Local   Local     00h08m10s  0
       system                                                      0
192.168.33.2/32                              Remote  BGP VPN   00h08m08s  170
       192.0.2.2 (tunneled)                                        0
--- snipped ---
===============================================================================
*A:LAC#
```

Operational PPPoE session information for IES 1 (base instance) is shown using following command.

```
*A:LAC# show service id 1 ppp session

===============================================================================
PPP sessions for service 1
===============================================================================
User-Name
  Descr.
           Up Time       Type  Termination     IP/L2TP-Id/Interface-Id MC-Stdby
-------------------------------------------------------------------------------
user1@wholesale.com
  svc:1 sap:1/1/3:100 mac:00:00:00:00:01:01 sid:1
           0d 00:01:20   oE    lac             1070684666
-------------------------------------------------------------------------------
No. of PPP sessions: 1
===============================================================================
*A:LAC#
```

Operational tunnel information for VPRN 65536 is displayed as follows.

```
*A:LAC# show router 65536 l2tp tunnel

===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State              Blacklist-state   Ses Active
  Group                                                             Ses Total
    Assignment
-------------------------------------------------------------------------------
1070661632 16337     14761     established        not-blacklisted   1
  RADIUS-returned-
TG                                                 1
    RADIUS-returned-TN
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
*A:LAC#
```

Operational session information for VPRN 65536 is displayed using following command, and shows that the session is established.

```
*A:LAC# show router 65536 l2tp session

===============================================================================
L2TP Session Summary
```

```
===============================================================================
ID                  Control Conn ID     Tunnel-ID   Session-ID  State
-------------------------------------------------------------------------------
1070684666          1070661632          16337       23034       established
-------------------------------------------------------------------------------
No. of sessions: 1
===============================================================================
*A:LAC#
```

Scenario 2: Node-Derived L2TP Parameters

RADIUS Returns L2TP Group

This example returns VPRN 65536 as the L2TP service instance [26-6527-104] Alc-Tunnel-Serv-Id and an l2tp group-name wholesale.com [26-6527-46] Alc-Tunnel-Group.

```
106 2016/06/03 13:54:59.17 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 4 len 122 from 172.16.1.1:1812 vrid 1 pol rsp-radius-1
    VSA [26] 21 Alcatel(6527)
      SUBSC ID STR [11] 19 user1@wholesale.com
    VSA [26] 15 Alcatel(6527)
      SUBSC PROF STR [12] 13 sub-profile-1
    VSA [26] 15 Alcatel(6527)
      SLA PROF STR [13] 13 sla-profile-1
    VSA [26] 6 Alcatel(6527)
      TUNNEL SERVICE ID [104] 4 65536
    VSA [26] 15 Alcatel(6527)
      TUNNEL GROUP [46] 13 wholesale.com
"
```

For operational PPPoE session information in IES 1/base instance, use following command.

```
*A:LAC# show service id 1 ppp session
===============================================================================
PPP sessions for service 1
===============================================================================
User-Name
  Descr.
           Up Time       Type  Termination     IP/L2TP-Id/Interface-Id MC-Stdby
-------------------------------------------------------------------------------
user1@wholesale.com
  svc:1 sap:1/1/3:100 mac:00:00:00:00:01:01 sid:1
           0d 00:01:53   oE    lac             233714323
-------------------------------------------------------------------------------
No. of PPP sessions: 1
===============================================================================
*A:LAC#
```

Operational tunnel information for VPRN 65536 shows the tunnel is in the established state.

```
*A:LAC# show router 65536 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State              Blacklist-state   Ses Active
  Group                                                             Ses Total
    Assignment
-------------------------------------------------------------------------------
233701376  3566      1245      established        not-blacklisted   1
  wholesale.com                                                     1

    wholesale.com
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
*A:LAC#
```

The operational session information for VPRN 65536 shows the session is in the established state.

```
*A:LAC# show router 65536 l2tp session
===============================================================================
L2TP Session Summary
===============================================================================
ID                  Control Conn ID     Tunnel-ID   Session-ID  State
-------------------------------------------------------------------------------
233714323           233701376           3566        12947       established
-------------------------------------------------------------------------------
No. of sessions: 1
===============================================================================
*A:LAC#
```

LUDB Returns L2TP Group

This example returns VPRN 65536 as the L2TP service instance and l2tp group-name wholesale.com (LUDB l2tp group "wholesale.com" service-id 65536).

The **debug subscriber-mgmt local-user-db "l2tp" detail all** command shows the LUDB authentication access (The returned parameter details are not shown).

```
139 2016/06/03 13:58:37.83 CEST MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  user-name:
    original:  user1@wholesale.com
    masked:    user1@wholesale.com

  Host wholesale.com found in user data base l2tp"
```

To show the operational data from LUDB *l2tp*, use the following command.

```
*A:LAC# show subscriber-mgmt local-user-db "l2tp" ppp-
host "wholesale.com" | match N/A invert-match | match none invert-match

===============================================================================
PPP Host "wholesale.com"
===============================================================================
Admin State          : Up
Last Mgmt Change     : 06/03/2016 13:43:14

Host Identification
 User Name           : wholesale.com (domain only)

Matched Objects      : userName

Password Type        : ignore
PADO Delay           : 0msec
Diameter app policy  : (Not Specified)
Diameter auth policy : (Not Specified)
Force IPv6CP         : Disabled
Ignore DF Bit        : Disabled

DHCPv6 lease times
 Renew timer         : > 9999 days
 Rebind timer        : > 9999 days
 Preferred lifetime  : 0d 00:00:00
 Valid lifetime      : 0d 00:00:00

Identification Strings (option 254)
 Subscriber Id       : user@wholesale.com
 SLA Profile String  : sla-profile-1
 Sub Profile String  : sub-profile-1

L2TP
 Service             : 65536
 Tunnel Group        : wholesale.com

MSAP defaults

Filter Overrules

Access loop info
===============================================================================
*A:LAC#
```

The **debug router l2tp event lac-session-setup** command shows the LAC tunnel selected for this example.

```
140 2016/06/03 13:58:37.83 CEST MINOR: DEBUG #2001 vprn65536 PPPoE 5->L2TP
"PPPoE 5->L2TP: UDP 192.168.33.1:1701 -> 192.168.33.2:1701
preference 50 tunnel wholesale.com:wholesale.com
    selected tunnel 3566"

141 2016/06/03 13:58:37.83 CEST MINOR: DEBUG #2001 vprn65536 PPPoE 5->L2TP
"PPPoE 5->L2TP: UDP 192.168.33.1:1701 -> 192.168.33.2:1701
preference 50 tunnel wholesale.com:wholesale.com
    create session 233723339"
```

For the operational PPPoE session information in IES 1/base instance, use the following command.

```
*A:LAC# show service id 1 ppp session

===============================================================================
PPP sessions for service 1
===============================================================================
User-Name
  Descr.
           Up Time       Type  Termination     IP/L2TP-Id/Interface-Id MC-Stdby
-------------------------------------------------------------------------------
user1@wholesale.com
  svc:1 sap:1/1/3:100 mac:00:00:00:00:01:01 sid:1
           0d 00:03:53   oE    lac             233723339
-------------------------------------------------------------------------------
No. of PPP sessions: 1
===============================================================================
*A:LAC#
```

Operational tunnel information for VPRN 65536 can be obtained using following command.

```
*A:LAC# show router 65536 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State              Blacklist-state   Ses Active
  Group                                                             Ses Total
    Assignment
-------------------------------------------------------------------------------
233701376  3566      1245      established        not-blacklisted   1
  wholesale.com                                                     1

    wholesale.com
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
*A:LAC#
```

The operational session information for VPRN 65536 shows the session is in the established state.

```
*A:LAC# show router 65536 l2tp session

===============================================================================
L2TP Session Summary
===============================================================================
ID                  Control Conn ID     Tunnel-ID   Session-ID  State
-------------------------------------------------------------------------------
233723339           233701376           3566        21963       established
-------------------------------------------------------------------------------
No. of sessions: 1
===============================================================================
*A:LAC#
```

Advanced Topics

Non-Session-Triggered L2TP Tunnel Setup

In addition to the ppp-session-triggered setup, an L2TP tunnel can also be set up via a tools command or an auto-establish command. These non-session-triggers are useful, for example, during the initial configuration phase where the LAC-LNS tunnel setup can be tested without any other user interaction. The PPP user still triggers the L2TP session-setup over this L2TP tunnel and RADIUS needs to return an l2tp group-name with the relevant name during authentication.

Auto-Establish

Every minute, a check is performed to determine if tunnels need to be established (a process referred to as scan auto-establish). The tunnel state is establishedIdle when the tunnel is setup, and becomes established when user triggered sessions are set up over this tunnel.

```
configure
    service
        vprn 65536 customer 1 create
            l2tp
                group "wholesale.com" create
                    tunnel "wholesale.com" create
                        local-address 192.168.33.1
                        local-name "lac-pe1"
                        peer 192.168.33.2
                        auto-establish
                        no shutdown
                    exit
                    no shutdown
                exit
                no shutdown
            exit
        exit
exit
```

There is no difference in operational behavior for a tunnel set up via a session-trigger or an auto-establish command. Removing the auto-establish parameter has no impact on active tunnels (establishedIdle or established).

```
*A:LAC# show router 65536 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State             Blacklist-state   Ses Active
  Group                                                            Ses Total
```

```
    Assignment
-------------------------------------------------------------------------------
1019609088 15558     3614      establishedIdle    not-blacklisted   0
  wholesale.com                                                     0
    wholesale.com
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
*A:LAC#
```

Tools Tunnel Start

An alternative for auto-establish is the tools start command.

```
configure
    service
        vprn 65536 customer 1 create
            l2tp
                group "wholesale.com" create
                    tunnel "wholesale.com" create
                        local-address 192.168.33.1
                        local-name "lac-pe1"
                        peer 192.168.33.2
                        no auto-establish
                        no shutdown
                    exit
                    no shutdown
                exit
                no shutdown
            exit
            no shutdown
        exit
    exit
exit

tools perform router 65536 l2tp group "wholesale.com" tunnel "wholesale.com" start
```

How Long Does a Tunnel Remain Idle Before Being Torn Down?

A persistent tunnel is a tunnel that remains available after the last session over that tunnel is closed. To create a persistent tunnel, the idle-timeout parameter must be set to infinite.

A non-persistent tunnel is torn down immediately (idle-timeout zero) after the last session over that tunnel is closed or after a configurable delay. The idle-timeout parameter is set via the RADIUS [26-6527-49] Alc-Tunnel-Idle-Timeout attribute or the corresponding node parameter. The default value for this parameter is infinite (persistent).

Idle-Timeout

```
configure router l2tp | configure service vprn l2tp
       idle-timeout [0..3600]seconds
       ---snipped---
       group <tunnel-group-name>
           idle-timeout [0..3600]secondzs | infinite
           ---snipped---
           tunnel  <tunnel-name>
              idle-timeout [0..3600]seconds | infinite
              ---snipped---
```

The following shows an example of a non-persistent tunnel (idle-timeout 30 seconds). The tunnel changes state from established to establishedIdle when the last session is terminated. Idle-timeout seconds later, the session changes to the closed state. For the purpose of troubleshooting, the operational data stays available for destruct-timeout seconds (see later).

```
*A:LAC# show router l2tp tunnel detail
===============================================================================
L2TP Tunnel Status
===============================================================================

Connection ID: 43646976
State        : closed

--- snipped ---

Name         : lac-pe1
Remote Name  : lns-pe2
Assignment ID: unnamed
Group Name   : default_radius_group
Acct. Policy : N/A
Error Message: idle timeout (30 seconds) expired

                                       Remote Conn ID    : 682491904
Tunnel ID         : 666                Remote Tunnel ID  : 10414
Preference        : 50                 Receive Window    : 64
Hello Interval (s): 300
Idle TO (s)       : 30                 Destruct TO (s)   : 60
Max Retr Estab    : 5                  Max Retr Not Estab: 5
Session Limit     : 32767              AVP Hiding        : never
Transport Type    : udpIp              Challenge         : never

--- snipped ---
```

```
No. of tunnels: 1
===============================================================================
*A:LAC#
```

The following shows an example of a persistent tunnel (idle-timeout infinite).

```
*A:LAC# show router l2tp tunnel detail
===============================================================================
L2TP Tunnel Status
===============================================================================

Connection ID: 209190912
State        : establishedIdle

--- snipped ---

Name         : lac-pe1
Remote Name  : lns-pe2
Assignment ID: unnamed
Group Name   : default_radius_group
Acct. Policy : N/A
Error Message: N/A

                                      Remote Conn ID    : 880410624
Tunnel ID         : 3192              Remote Tunnel ID  : 13434
Preference        : 50                Receive Window    : 64
Hello Interval (s): 300
Idle TO (s)       : infinite          Destruct TO (s)   : 60
Max Retr Estab    : 5                 Max Retr Not Estab: 5
Session Limit     : 32767             AVP Hiding        : never
Transport Type    : udpIp             Challenge         : never

--- snipped ---

No. of tunnels: 1
===============================================================================
*A:LAC#
```

Tools Tunnel Stop

In addition to the idle-timeout used for tunnel termination, a tools stop command is also available that can be used to terminate persistent and non-persistent tunnels at any moment in time. Be aware that this command is very destructive and destroys all sessions carried over the closed tunnel.

Following command shows the tunnel is in the establishedIdle state.

```
*A:LAC# show router 65536 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State             Blacklist-state   Ses Active
  Group                                                            Ses Total
    Assignment
-------------------------------------------------------------------------------
```

```
246415360  3760      2730      establishedIdle    not-blacklisted   0
  wholesale.com                                                    1
    wholesale.com
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
*A:LAC#
```

The command below terminates the l2tp tunnel. The tunnel is aborted (the LAC sends StopCCN) using the <connection-id> or <tunnel-group-name>+<tunnel-name> as input. This StopCCN indicates "operator request" as the error reason.

```
*A:LAC# tools perform router 65536 l2tp group "wholesale.com" tunnel "wholesale.com"
 stop
INFO: CLI stopped 1 tunnels, destructed 0 tunnels.
```

The following debug output shows the tunnel being aborted.

```
90 2016/06/06 08:52:30.56 CEST MINOR: DEBUG #2001 vprn65536 L2TP(v2, ctrl, egres
s)
"L2TP(v2, ctrl, egress): UDP 192.168.33.1:1701 -> 192.168.33.2:1701
tunnel 2730 session 0, ns 5 nr 2, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        StopControlConnectionNotification(4)
    AVP ResultCode(0,1), flags: mandatory, reserved=0
        result-code: "generalRequestToClearControlConnection"(1), error-code: "n
oGeneralError"(0)
        error-msg: "operator request"
    AVP AssignedTunnelId(0,9), flags: mandatory, reserved=0
        3760"
```

Alternatively, the tunnel can also be stopped with the following command. The effect would be the same.

```
*A:LAC# tools perform router 65536 l2tp tunnel 246415360 stop
```

Keepalive (hello)

A keepalive mechanism is employed by L2TP in order to differentiate tunnel outages from extended periods of no control or data activity on a tunnel. This is accomplished by injecting Hello control messages after a specified period of time has elapsed since the last data or control message (ZLB not included) was received on a tunnel. As for any other L2TP control message, if the Hello message is not reliably delivered, then the tunnel is declared down and reset, as defined in RFC 2661, *Layer Two Tunneling Protocol "L2TP"*. This means that SROS does not initiate hello packets if session control traffic is handled over this tunnel. The hello timer is reset if the system transmits any control packet over this tunnel (ZLB packets and data traffic are not taken into account).

The keepalive function is disabled (not recommended) using RADIUS [26-6527-50] Alc-Tunnel-Hello-Interval -1 or hello-interval infinite (default 300). The number of retries for unsuccessful Hello packet delivery equals RADIUS [26-6527-52] Alc-Tunnel-Max-Retries-Estab or node parameter max-retries-estab (default 5). The retry interval is initially set to 1 second and doubles on each retry with a maximum interval of 8 seconds. Using a max-retries-estab 7 results in a retry of [1,2,4,8,8,8,8 seconds].

```
configure router l2tp | configure service vprn l2tp
    hello-interval [60..3600]seconds | infinite
    max-retries-estab [2..7]
    ---snipped---
    group <tunnel-group-name>
        hello-interval [60..3600]seconds | infinite
        max-retries-estab [2..7]
        ---snipped---
        tunnel  <tunnel-name>
            hello-interval [60..3600]seconds | infinite
            max-retries-estab [2..7]
            ---snipped---
```

For example, the LAC can be configured with an hello-timer of 1 minute and the LNS with an hello-timer of 2 minutes. The hello-timer interval for LAC and LNS do not have to be same as the keepalive mechanism works asynchronous. See L2TP Keepalive Mechanism.

```
*A:LAC# show router l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State              Blacklist-state   Ses Active
  Group                                                             Ses Total
    Assignment
-------------------------------------------------------------------------------
744751104  11364     5391      established        not-blacklisted   1
  default_radius_group                                              1
    unnamed
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
*A:LAC#
```

Figure 179 **L2TP Keepalive Mechanism**

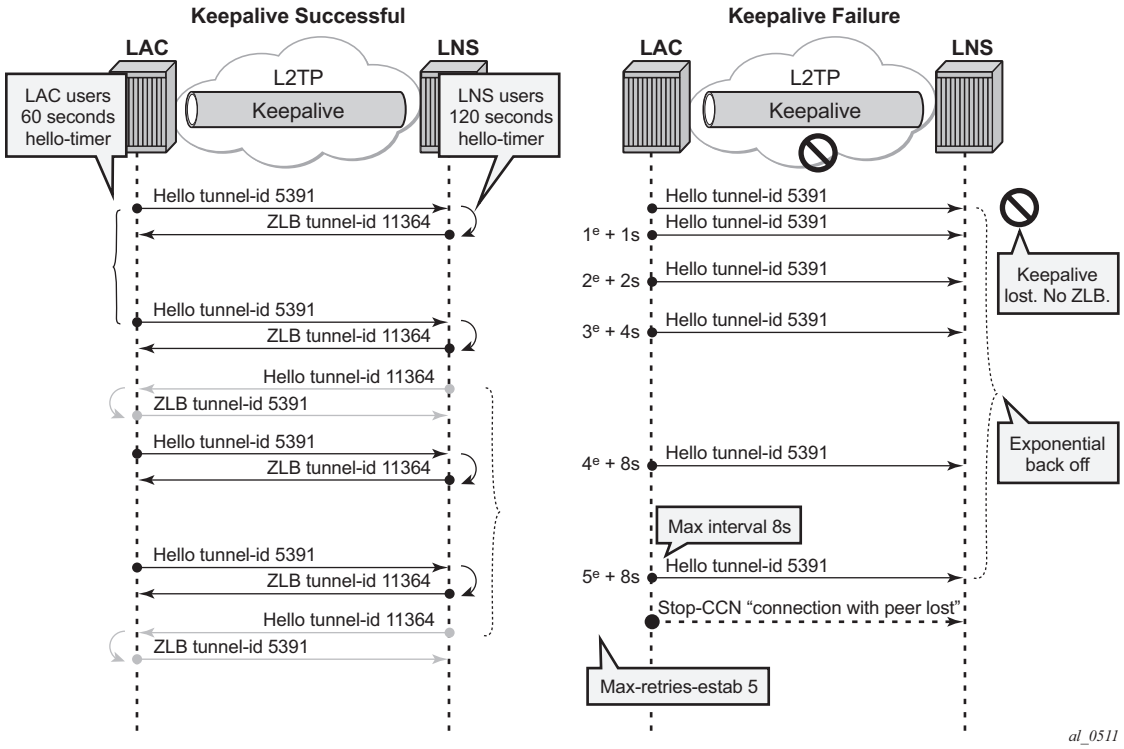

Figure 179 shows the tunnel being closed after 5 unsuccessful Hello deliveries with error-message **connection with peer lost**.

```
*A:LAC# show router l2tp tunnel detail
===============================================================================
L2TP Tunnel Status
===============================================================================
Connection ID: 744751104
State        : closed
IP           : 192.0.2.1
UDP          : 1701
Peer IP      : 192.168.22.2
Peer UDP     : 1701
Tx dst-IP    : 192.168.22.2
Tx dst-UDP   : 1701
Rx src-IP    : 192.168.22.2
Rx src-UDP   : 1701
Name         : lac-pe1
Remote Name  : lns-pe2
Assignment ID: unnamed
Group Name   : default_radius_group
Acct. Policy : N/A
Error Message: connection with peer lost
                                        Remote Conn ID    : 353304576
Tunnel ID         : 11364               Remote Tunnel ID  : 5391
Preference        : 50                  Receive Window    : 64
Hello Interval (s): 300
```

```
--- snipped ---

No. of tunnels: 1
===============================================================================
*A:LAC#
```

Does L2TP Keep Info About Closed Tunnels, Sessions?

The destruct-timeout parameter (expressed in seconds) controls the period of time that the tunnel, or session data related to a closed (disconnected) tunnel, or session persists before being removed. The destruct_timeout is a debugging aid by saving underlying memory structures after the tunnel, or session is terminated. It is configured via the RADIUS [26-6527-51] Alc-Tunnel-Destruct-Timeout attribute or the corresponding node parameter. Default value for this parameter is 60 seconds.

```
configure router l2tp | configure service vprn l2tp
       destruct-timeout [60..86400]
       ---snipped---
       group <tunnel-group-name>
           destruct-timeout [60..86400]
           ---snipped---
           tunnel  <tunnel-name>
              destruct-timeout [60..86400]
```

The following output shows a session that is closed and the reason for it being terminated.

```
*A:LAC# show router l2tp session detail
===============================================================================
L2TP Session 833297393
===============================================================================
Connection ID: 833297393
State        : closed
Tunnel Group : default_radius_group
Assignment ID: unnamed
Error Message: Terminated by PPPoE: Received PPPoE PADT
Control Conn ID   : 833290240           Remote Conn ID    : 590235074
Tunnel ID         : 12715               Remote Tunnel ID  : 9006
Session ID        : 7153                Remote Session ID : 17858
Time Started      : 06/06/2016 09:21:30
Time Established  : 06/06/2016 09:21:30 Time Closed       : 06/06/2016 09:25:25
CDN Result        : generalError        General Error     : vendorSpecific
-------------------------------------------------------------------------------
No. of sessions: 1
===============================================================================
*A:LAC#
```

The following output shows a tunnel that is closed and the reason for it being closed.

```
*A:LAC# show router l2tp tunnel detail
===============================================================================
L2TP Tunnel Status
===============================================================================

Connection ID: 833290240
State        : establishedIdle
IP           : 192.0.2.1
UDP          : 1701
Peer IP      : 192.168.22.2
Peer UDP     : 1701
Tx dst-IP    : 192.168.22.2
Tx dst-UDP   : 1701
Rx src-IP    : 192.168.22.2
Rx src-UDP   : 1701
Name         : lac-pe1
Remote Name  : lns-pe2
Assignment ID: unnamed
Group Name   : default_radius_group
Acct. Policy : N/A
Error Message: N/A

                                       Remote Conn ID    : 590217216
Tunnel ID         : 12715              Remote Tunnel ID  : 9006
Preference        : 50                 Receive Window    : 64
Hello Interval (s): 300
Idle TO (s)       : infinite           Destruct TO (s)   : 60
Max Retr Estab    : 5                  Max Retr Not Estab: 5
Session Limit     : 32767              AVP Hiding        : never
Transport Type    : udpIp              Challenge         : never
Time Started      : 06/06/2016 09:21:30 Time Idle         : 06/06/2016 09:25:25
Time Established  : 06/06/2016 09:21:30 Time Closed       : N/A
Stop CCN Result   : noError            General Error     : noError
Blacklist-state   : not-blacklisted
Set Dont Fragment : true

Failover
State             : not-recoverable
Recovery Conn ID  : N/A
Recovery state    : not-applicable
Recovered Conn ID : N/A
Recovery method   : mcs
Track SRRP        : (Not specified)
Ctrl msg behavior : handle
Recovery time (ms)
Requested         : N/A
Peer              : N/A
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
*A:LAC#
```

Floating Peers

A floating peer exists if the peer LNS address indicated in the source address of the SCCRP is different from the peer address known on the LAC. Floating peer allowance is configuration driven and is rejected by default.

The parameter peer-address-change-policy specifies whether the LAC accepts, ignores or rejects requests from a peer to change the destination IP address or UDP port.

```
configure router l2tp | configure service vprn l2tp
       peer-address-change-policy accept | ignore | reject
```

- accept — Specifies that this system accepts any source IP address change for received L2TP control messages related to a locally originated tunnel in the state wait-reply and rejects any peer address change for other tunnels. In case the new peer IP address is accepted, it is learned and used as destination address in subsequent L2TP messages.

Figure 180 **Floating Peers Accept**

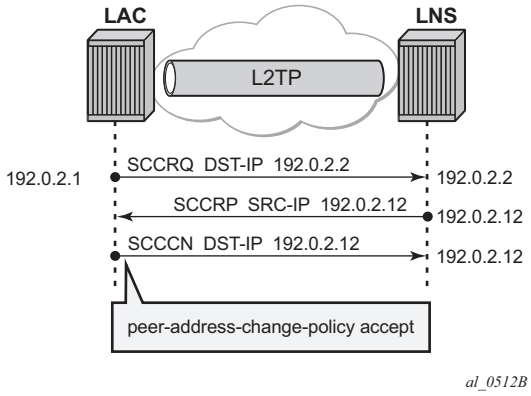

- Ignore — Specifies that this system ignores any source IP address change for received L2TP control messages, does not learn any new peer IP address and does not change the destination address in subsequent L2TP messages.

Figure 181 **Floating Peers Ignore**

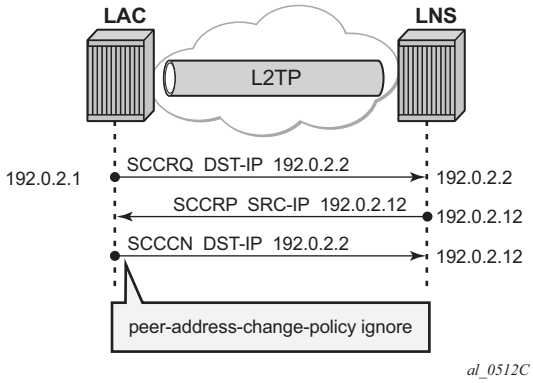

- Reject — Specifies that this system rejects any source IP address change for received L2TP control messages and drops those messages.

Figure 182 **Floating Peers Reject**

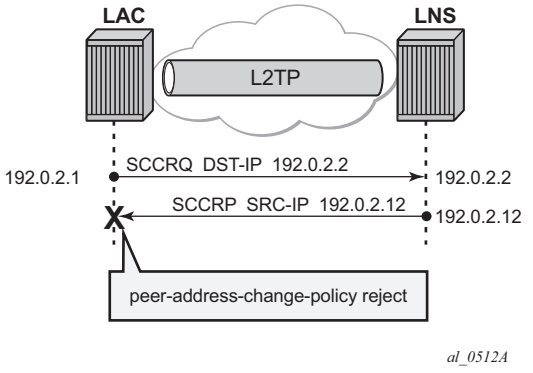

The values Peer IP, Tx dst-IP and Rx src-IP in the **show router l2tp tunnel detail** command indicates if floating peers are used or not.

An example of a floating peer (peer-address-change-policy accept) is as follows.

```
*A:LAC# show router l2tp tunnel detail
=====================================
L2TP Tunnel Status
=====================================
Connection ID: 897122304
State        : established
IP           : 192.0.2.1
UDP          : 1701
Peer IP      : 192.0.2.2 # (1) peer address used in SCCRQ
Peer UDP     : 1701
Tx dst-IP    : 192.0.2.12 # (3) peer address used in SCCCN
```

```
Tx dst-UDP   : 1701
Rx src-IP    : 192.0.2.12 # (2) SCCRP different IP received
Rx src-UDP   : 1701
---snipped---
```

Tx Connect Speed AVP 24 and Rx Connect Speed AVP 38

The Connect Speed (TX AVP 24 and RX AVP 38) is passed in the ICCN messages sent from the LAC to the LNS. The L2TP AVP 24 defines the (Tx) connect speed in bps from the perspective of traffic flowing from the LAC towards the subscriber (BNG downstream rate).The L2TP AVP 38 defines the (Rx) connect speed in bps from the perspective of traffic flowing from the subscriber towards the LAC (BNG upstream rate).

The report-rate configuration option indicates what rate is reported to the LNS when creating an L2TP session.

```
configure subscriber-mgmt
 sla-profile <sla-profile-name>
  ingress
   report-rate agg-rate-limit|scheduler|pppoe-actual-rate|rfc5515-actual-rate
  egress
   report-rate agg-rate-limit|scheduler|pppoe-actual-rate|rfc5515-actual-rate
```

- agg-rate-limit — Take the aggregate rate as received from the RADIUS Access-Accept message in VSA Alc-Subscriber-QoS-Override. When this RADIUS VSA is not present in the Access-Accept, or when RADIUS is not used, then take the configured aggregate rate limit. In the case where this is not configured, then take the port rate.
- scheduler <scheduler-name> — Take the rate of the specified scheduler. In case the scheduler is not linked with the scheduler-policy from the subscriber-profile, then take the port rate.
- pppoe-actual-rate — Take the rate from the DSL-Forum Vendor-Specific PPPoE Tag when available, otherwise take the port rate.
- rfc5515-actual-rate — Put 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.

Calling Number AVP 22 Format

The format of AVP 22 Calling Number in the ICRQ message is configurable via the parameter calling-number-format. The default format is "%S<space>%s" and corresponds to the concatenation of system-name<space>sap-id. Available parameters are %S (system-name), %c (Agent Circuit Id), %r Agent Remote Id, %s (sap-id), %l (Logical Line ID) and fixed strings. A combination can be configured from any of these parameters, but the total configured format cannot exceed 255 characters.

Example 1: Default configuration.

```
configure router l2tp calling-number-format "%S %s"


131379 2016/06/06 09:43:52.27 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress
)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 9006 session 0, ns 8 nr 6, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallRequest(10)
    AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
        21406
    AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
        22473
    AVP CallingNumber(0,22), flags: mandatory, reserved=0
        "LAC 1/1/3:100""
```

Example 2: Customized configuration and all parameters (%S %s %c) are available to construct the requested AVP 22.

```
configure router l2tp calling-number-format "start-%S###%s###%c-end"


131407 2016/06/06 09:52:12.98 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress
)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 9006 session 0, ns 12 nr 8, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallRequest(10)
    AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
        3949
    AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
        22474
    AVP CallingNumber(0,22), flags: mandatory, reserved=0
        "start-LAC###1/1/3:100###-end""
```

Example 3: Customized configuration and not all parameters are available to construct the requested AVP 22. Option-82 circuit-id (%c),remote-id (%r), and LLID (%l) information are lacking and therefore missing (skipped) in the formatted attribute.

```
configure router l2tp calling-number-format "%S#%c#%r#%l#%s"
```

```
131431 2016/06/06 09:54:46.68 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress
)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 9006 session 0, ns 15 nr 9, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallRequest(10)
    AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
        30074
    AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
        22475
    AVP CallingNumber(0,22), flags: mandatory, reserved=0
        "LAC####1/1/3:100""
```

Prevent LAC from Transmitting Calling Number AVP 22 to LNS

By default, the LAC includes the Calling Number AVP 22 in the L2TP incoming-call-request (ICRQ) packets transmitted to LNS. This AVP identifies the interface that is connected to the customer in the access network. Network access interface information can be hidden by configuring the LAC not to send the Calling Number AVP to the LNS.

Use the following command to disable the sending of L2TP Calling Number AVP 22.

```
configure router l2tp
        exclude-avps calling-number
```

cisco-nas-port AVP 100

Interoperation with a Cisco LNS requires that the LAC communicates a NAS port type to the LNS via the L2TP ICRQ ‘Cisco Nas Port Info AVP (100)’. This AVP (100) includes information that identifies the NAS port and indicates whether the port type is Ethernet or ATM and is configured via the cisco-nas-port parameter.

Cisco AVP 100 format

- First 5 bytes are NAS-Port-Type:
  - 0f10090203 (Ethernet)
  - 0f10090201 (ATM)
- Remaining 4 bytes corresponds with the configured cisco-nas-port value

Example:

- Ethernet 12b s-vlan-id; 10b c-vlan-id; 3b slot number; 2b MDA nbr; 5b port

- ATM 12b VPI; 10b VCI; 3b slot number; 2b MDA nbr; 5b port

```
configure router l2tp
        cisco-nas-port ethernet "*12o*10i*3s*2m*5p" atm "*12v*10c*3s*2m*5p"
```

nas-port 1/1/3:100 corresponds to 102563 (000000000000 0001100100 001 01 00011).

```
131455 2016/06/06 09:57:38.57 CEST MINOR: DEBUG #2001 Base L2TP(v2, ctrl, egress
)
"L2TP(v2, ctrl, egress): UDP 192.0.2.1:1701 -> 192.168.22.2:1701
tunnel 9006 session 0, ns 18 nr 10, flags:, reserved=0
    AVP MessageType(0,0), flags: mandatory, reserved=0
        IncomingCallRequest(10)
    AVP CiscoNasPort(9,100), flags:, reserved=0
        102563 type=ethernet(0f:10:09:02:03)
    AVP AssignedSessionId(0,14), flags: mandatory, reserved=0
        26023
    AVP CallSerialNumber(0,15), flags: mandatory, reserved=0
        22476
    AVP CallingNumber(0,22), flags: mandatory, reserved=0
        "LAC####1/1/3:100""
```

L2TP Group/Peer/Tunnel Draining

When the LAC has established sessions, the LAC can avoid the creation of new sessions for a specific group, peer, or tunnel, via the **drain** command.

No new sessions are created for a group, peer or tunnel that is being drained (draining state) but the current sessions are left intact.

After the **drain** command is issued, the group, peer, or tunnel moves from a draining to drained state when the last session is closed. A drained group, peer, or tunnel can then be managed (reconfigured, deleted) without any user impact.

Be aware that a group, peer, or tunnel in a draining or drained state is skipped in the tunnel selection process. The next example shows a tunnel draining; group and peer draining works according in the same way.

A tunnel has 1 session and is in established state.

```
*A:LAC# show router 65536 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State              Blacklist-state   Ses Active
  Group                                                              Ses Total
    Assignment
-------------------------------------------------------------------------------
25886720   395       11602     established        not-blacklisted   1
  wholesale.com                                                      1
    wholesale.com
-------------------------------------------------------------------------------
```

```
No. of tunnels: 1
===============================================================================
*A:LAC#
```

The following tools **drain** command puts the tunnel in a draining state and leaves the sessions intact.

```
*A:LAC# tools perform router 65536 l2tp tunnel 25886720 drain
```

Initially the tunnel is in the draining state.

```
*A:LAC# show router 65536 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State             Blacklist-state   Ses Active
  Group                                                            Ses Total
    Assignment
-------------------------------------------------------------------------------
25886720   395       11602     draining          not-blacklisted   1
  wholesale.com                                                    1
    wholesale.com
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
*A:LAC#
```

The tunnel moves to the drained state at the moment the last session is closed. Debugging shows that a drained tunnel is also not used as last resort and is skipped during the tunnel selection process.

```
*A:LAC# show router 65536 l2tp tunnel
===============================================================================
Conn ID    Loc-Tu-ID Rem-Tu-ID State             Blacklist-state   Ses Active
  Group                                                            Ses Total
    Assignment
-------------------------------------------------------------------------------
25886720   395       11602     drained           not-blacklisted   0
  wholesale.com                                                    1
    wholesale.com
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
*A:LAC#
```

The output below shows new sessions cannot select a drained tunnel.

```
154718 2016/06/06 10:54:13.73 CEST MINOR: DEBUG #2001 vprn65536 PPPoE 1420->L2TP
"PPPoE 1420->L2TP: UDP 192.168.33.1:1701 -> 192.168.33.2:1701
preference 50 tunnel wholesale.com:wholesale.com
    no additional session can be created in tunnel 14567"

154719 2016/06/06 10:54:13.73 CEST MINOR: DEBUG #2001 vprn65536 PPPoE 1420->L2TP
"PPPoE 1420->L2TP:
    stop: no more tunnels can be tried"
```

The drained tunnel can then be closed without user impact.

```
tools perform router 65536 l2tp tunnel 954662912 stop
```

For draining and undraining for example a group, following commands can be used.

```
tools perform router 65536 l2tp group "wholesale.com" drain
tools perform router 65536 l2tp group "wholesale.com" no drain
```

Conclusion

This example provides the LAC L2TP access server configuration and troubleshooting commands for the LAA architecture (tunneled-access) model.

Local User Database Basics

This chapter provides information about Local User Database (LUBD) Basics.

Topics in this chapter include:

Applicability

This note is applicable to the 7x50 SR series and was tested on SR-OS 13.0.R1. Chassis mode B or later must be used.

Overview

A Local User Database (LUDB) is a data source containing a set of host entries, providing full or partial Enhanced Subscriber Management (ESM) data so that subscribers and subscriber hosts can be instantiated when end-users connect their devices.

An LUDB can be accessed for the following applications; see Figure 183.

1. To support ESM for retrieval of data to instantiate hosts and subscribers. This applies to the Routed Central Office [CO] model only.
2. To support a local DHCPv4 server; for example for assigning fixed IP addresses to dedicated end-user devices.
3. To allow the system to provide the ESM data in case the RADIUS server referenced from the authentication policy is not available. The LUDB serves as a fallback for RADIUS authentication.

Figure 183 **LUDB Applications**

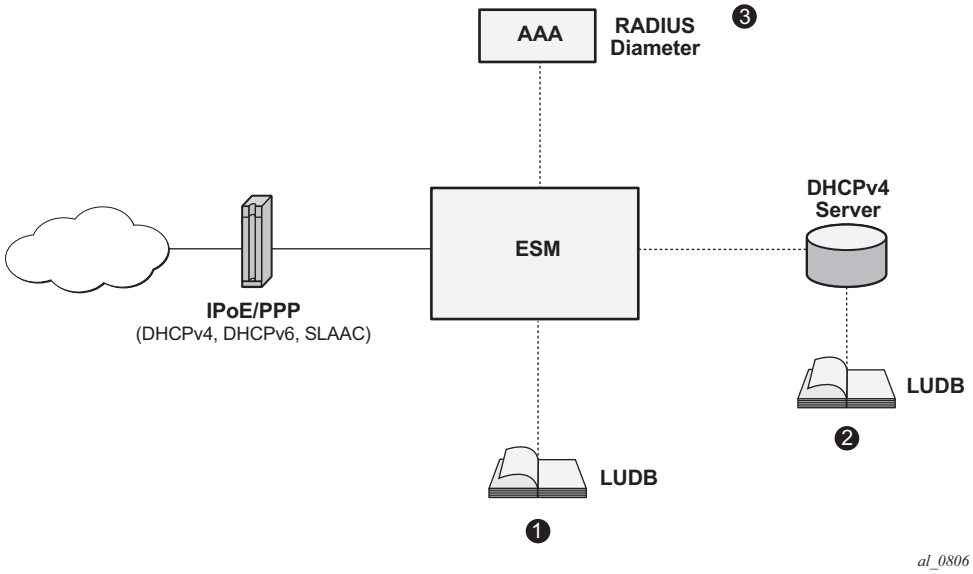

The LUDB lookup process is common to the applications shown in Figure 183, and performs the following steps:

- Applying match criteria, to select the input parameters that will be used for the lookup
- Optionally, applying a mask to one or more of the remaining input parameters
- Performing the lookup

The LUDB lookup process translates a set of input parameters (the host identification fields) into a set of output parameters; see Figure 184 for the following example:

- An LUDB lookup is requested for a client with MAC address and SAP as input parameters (1).
- The match criteria indicate to consider the SAP only, so the MAC address is ignored (2).
- The masking defines the stripping of the VLAN-tag from the SAP (3).
- The lookup then uses SAP 1/1/1 and finds entry1 to be the matching entry, so the LUDB returns the SLA-profile string and the SUB-profile string together with the Gi address.

Optionally, an LUDB defines a **default** host entry, which is used in case none of the other entries matches the lookup request, so it serves as a wildcard (*).

Not finding any matching host entry in an LUDB results in a setup failure.

Figure 184 **Processing an LUDB Lookup Request**

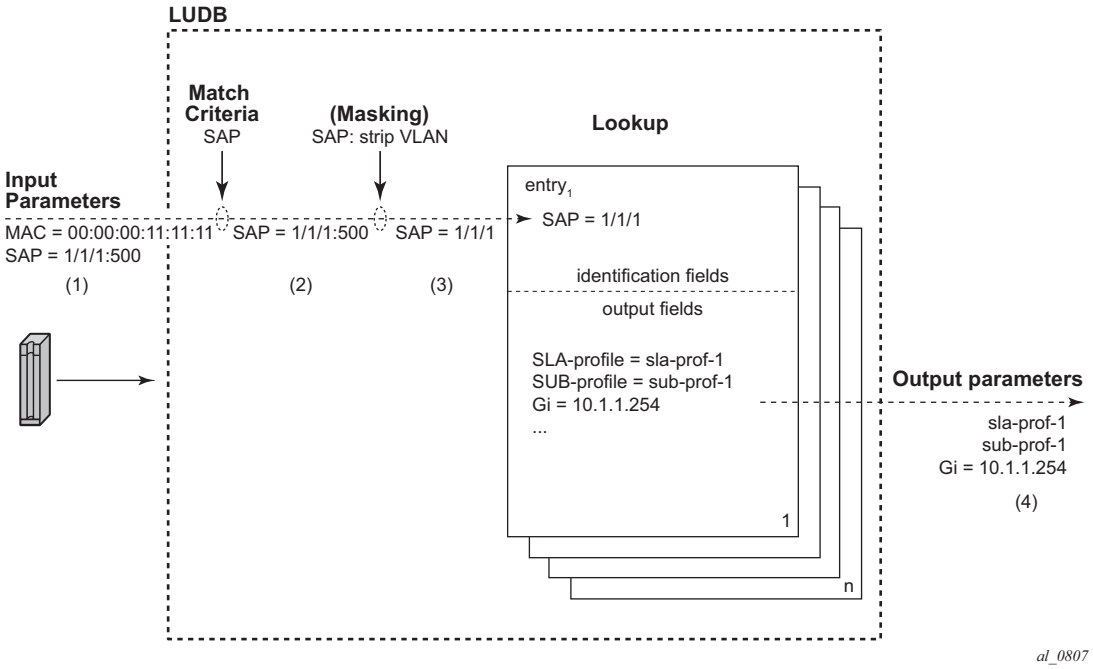

Configuration

Creating LUDBs

An LUDB is identified by a name of 32 characters maximum (Figure 185).

```
*A:BNG-1>config>subscr-mgmt# local-user-db
  - local-user-db <local-user-db-name> [create]
```

Multiple LUDBs can be defined, and their respective names must be unique.

An LUDB can provide the data for IPoE (DHCPv4, DHCPv6, and SLAAC) as well as for PPP users.

```
*A:BNG-1>config>subscr-mgmt>loc-user-db$
 [no] description     - Description for this local user database
      ipoe            + Configure IPOE hosts
      ppp             + Configure PPP hosts
 [no] shutdown        - Administratively enable/disable this local user database
```

For an LUDB to be active, the LUDB must be in the **no shutdown** state.

Individual host entries in an LUDB can match single or multiple hosts.

Figure 185 **Creating LUDBs and LUDB Entries**

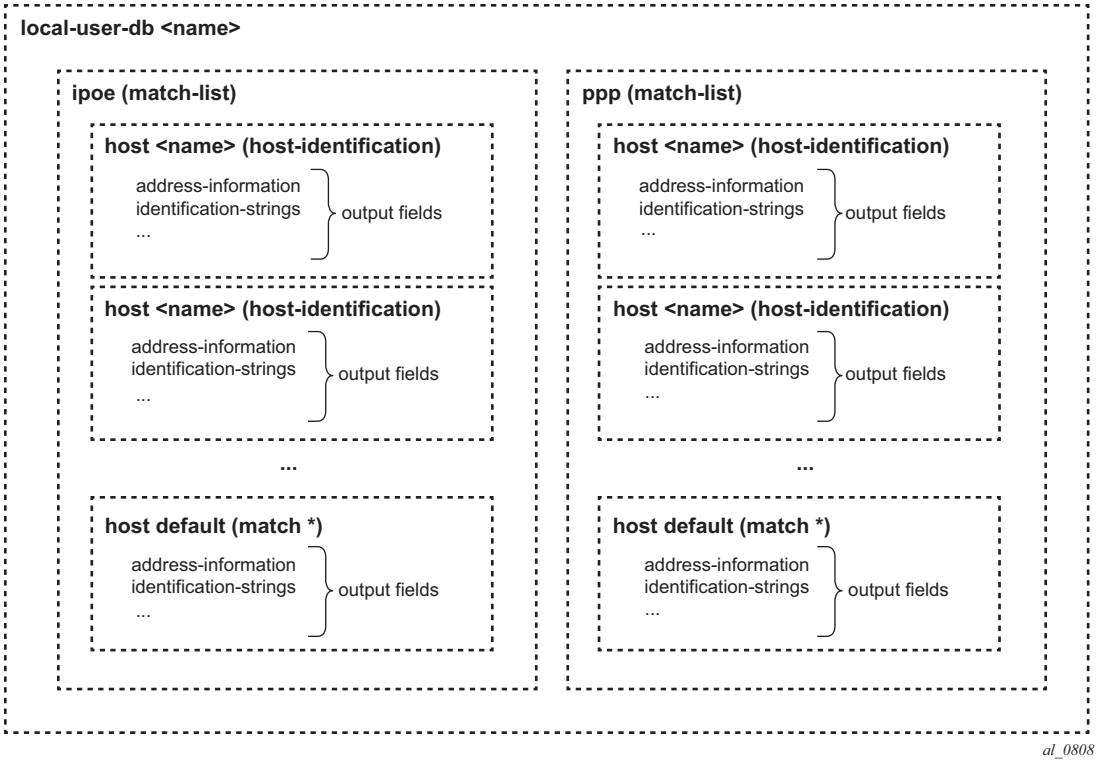

Creating Host Entries

A host entry is identified by name of 32 characters maximum (Figure 185).

```
*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe$ host
  - host <host-name> [create]
```

The name **default** can optionally be used as a wildcard for situations where otherwise the lookup fails to find a matching entry. A host entry belongs to either the IPoE or the PPP section of an LUDB. The name of a host entry must be unique within the section. A host entry contains two sets of fields. The first set of fields are the host-identification fields and are used for the lookup, the second set of fields are output to the lookup process.

The host-identification fields available for IPoE are, in alphabetical order:

- circuit-id
- derived-id, which must be defined using a Python script, which derives the value from DHCP messages
- encap-tag-range
- mac-address
- option 60
- remote-id
- sap-id
- service-id
- string
- system-id

The host-identification fields available for PPP are, in alphabetical order:

- circuit-id: taken from the PPPoE tags
- encap-tag-range
- mac
- remote-id: taken from the PPPoE tags
- sap-id
- service-name
- username

The output fields of the lookup process include the identification strings, DHCP options, IP address information, MSAP information, and so on.

Entry Validation

For a host entry to be active, it must be put in the **no shutdown** state.

Before adding the host entry to the lookup database, the system validates the host entry:

- A **default** host entry can be added, preferably without host identification fields.

```
*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe# host default create
*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe>host$ no shutdown
  INFO: DHCPS #1138 This host will be considered as the default host
*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe>host$
```

Defining a **default** host entry with identification fields is not recommended, because it would turn the default host entry into a regular entry, instead of a match all entry when the match-list is changed.

- A **non default** host entry without host identification fields cannot be added to the lookup database.

```
MINOR: DHCPS #1126 Host-identification must have at least 1 item defined
```

- A **non default** host entry with none of its host identification fields in common with the match-list is added to the unmatched host list.

```
INFO: DHCPS #1107 Host could not be inserted in lookup database -
no match values
```

- A **non default host** entry is added to the lookup database when at least one of the defined host identification fields is common with the match-list, even when some of the host identification fields are not on the match-list.
- Two or more **non default** host entries with the same host-identification definitions are considered as duplicates. The second entry with the same host-identification definitions is not added to the lookup database; it is considered as a configuration mistake.

```
INFO: DHCPS #1107 Host could not be inserted in lookup database — duplicate
```

Note that LUDB informational and error messages appear to be originating from the DHCPS application (DHCPS #nnn in the preceding outputs), even though the LUDB is not associated with a DHCPv4 server.

Creating a Match-List

Retrieving data from an LUDB requires one or more criteria to be put on a **match-list**. A match-list is a sequential list of parameters considered for the lookup; other parameters provided on LUDB access are ignored.

For IPoE, up to four criteria can be defined; for PPP, the maximum is three. The criteria on a match-list are processed in the order specified.

For IPoE users, the following match criteria are allowed, in alphabetical order:

- circuit-id
- derived-id (defined by a Python script)
- dual-stack-remote-id (IPv4 and IPv6, with IPv6 enterprise-id stripped off)
- encap-tag-range
- mac-address

- option 60
- remote-id (IPv4 and IPv6, including the IPv6 enterprise-id)
- sap-id
- service-id
- string
- system-id

```
A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe# match-list
  - no match-list
  - match-list <ipoe-match-type-1> [<ipoe-match-type-2>...(up to 4 max)]

 <ipoe-match-type>    : circuit-id|derived-id|dual-stack-remote-id|encap-tag-
range|mac|option60|remote-id|sap-id|service-id|string|system-id

A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe#
```

For PPP users, the following match criteria are allowed, in alphabetical order:

- circuit-id
- encap-tag-range
- mac-address
- remote-id (IPv4 and IPv6, including the IPv6 enterprise-id)
- sap-id
- service-name
- username (complete username, domain part only, host part only)

```
A:BNG-1>config>subscr-mgmt>loc-user-db>ppp# match-list
- no match-list
- match-list <ppp-match-type-1> [<ppp-match-type-2>...(up to 3 max)]

<ppp-match-type>     : circuit-id|mac|remote-id|sap-id|encap-tag-range|service-
name|username

A:BNG-1>config>subscr-mgmt>loc-user-db>ppp#
```

Masking

Optionally, the parameters considered for the lookup can be masked.

Masking is prefix- or suffix- based, or a combination of both. A prefix or suffix string, or a prefix or suffix length, can be specified.

For PPP users, masks can be applied to the circuit-id, remote-id, sap-id, service-name, and username. For IPoE users, masks can be applied to the circuit-id, option 60, remote-id, sap-id, string, and system-id.

```
*A:BNG-1>config>subscr-mgmt>loc-user-db>ppp# mask
  - 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>

 <ppp-match-type>     : circuit-id|remote-id|sap-id|service-name|username
 <prefix-string>      : [127 chars max] ('*' is wildcard)
 <prefix-length>      : [1..127]
 <suffix-string>      : [127 chars max] ('*' is wildcard)
 <suffix-length>      : [1..127]

*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe# mask
  - 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>

 <ipoe-match-type>    : circuit-id|option60|remote-id|sap-id|string|system-id
 <prefix-string>      : [127 chars max] ('*' is wildcard)
 <prefix-length>      : [1..127]
 <suffix-string>      : [127 chars max] ('*' is wildcard)
 <suffix-length>      : [1..127]
```

The lookup occurs after applying the optional masks.

The examples in Table 40 illustrate masking. Note that for the third example, a combination of both prefix and suffix matching is used.

Table 40 **Masking Examples**

Mask Type			prefix-length	suffix-length	prefix-string	suffix-string	result
username	**circuit-id**	**remote-id**					
-	-	87654321-BSAN-1	9	-	-	-	BSAN-1
-	BSAN-2|1|100|1/2/1	-	-	11	-	-	BSAN-2
all@domain-1.com	-	-	-	-	*@	.com	domain-1

Lookup

The following rules apply while scanning through an LUDB in search of a single matching entry:

1. Only criteria on the match-list are considered.

   Assume a client with MAC-address, a circuit-id, and a remote-id. If the match-list only defines the MAC-address to be used as criterion, then the circuit-id and the remote-id are ignored. Only the MAC-address is used for selecting the proper host entry.
2. The order of the criteria on the match-list is important.

   The match-list is a sequential list, and the criteria are processed left to right.
3. As many of the host-identification fields as possible must be matched, while still obeying rule 1.

   Only the (optionally masked) parameters from the match list are verified.
4. A **default host** is excluded from the scan, if defined.

   A **default host** is used as a fallback when scanning through an LUDB does not provide any result.

Figure 186 **Host Matching Examples**

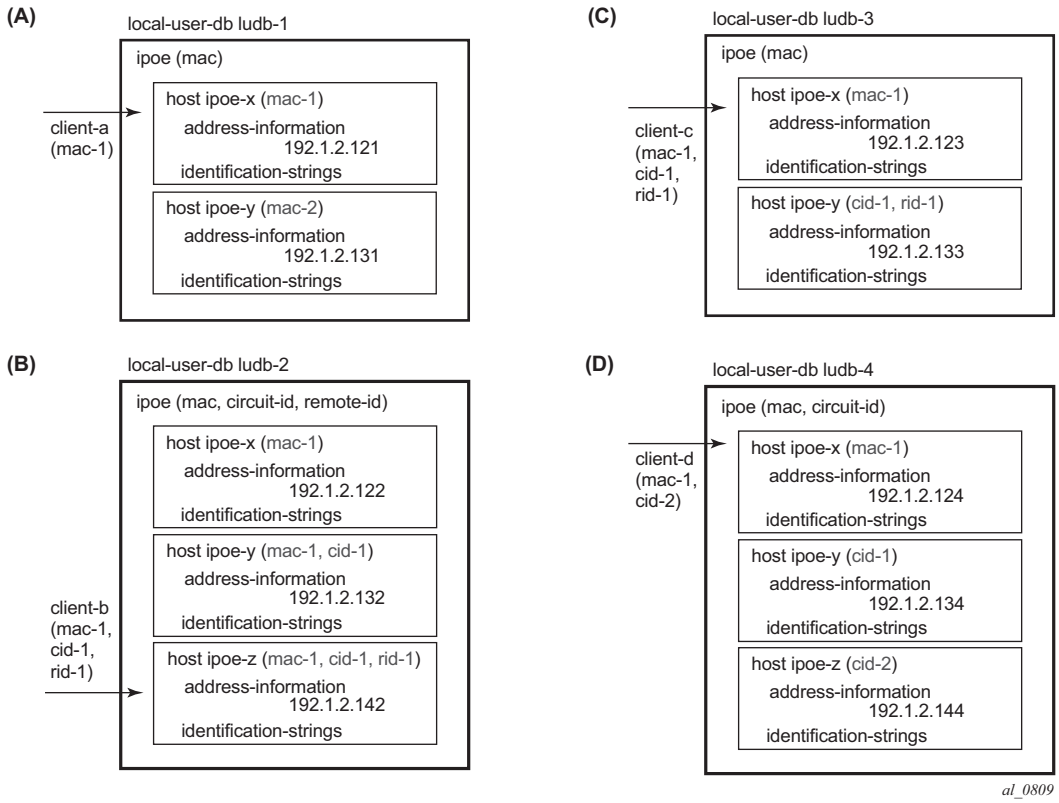

The examples in Figure 186 and Figure 187 illustrate these rules:

a. Matching is based on the MAC-address only. When **client-a** with **mac-1** connects, host **ipoe-x** is matched.

b. Matching is based on the MAC-address, circuit-id, and remote-id, in this sequence. As client-b enters with mac-1, cid-1 and rid-1, the match-list is scanned and matched left to right, so host ipoe-z is matched.

c. Matching is based on the MAC-address only. Even though client-c connects with mac-1, cid-1, and rid-1, the system ignores the circuit-id and the remote-id, so the matching host is ipoe-x. Note that host ipoe-y can never be matched using the match-list defined; it is on the unmatched host list.

d. Matching is based on the MAC-address and the circuit-id, in this sequence. Client-d connects with mac-1 and cid-2, but because the system scans the match-list left through right, the MAC-address takes priority over the circuit-id. The matching host is ipoe-x.

Figure 187 **Host Matching Examples (Continued)**

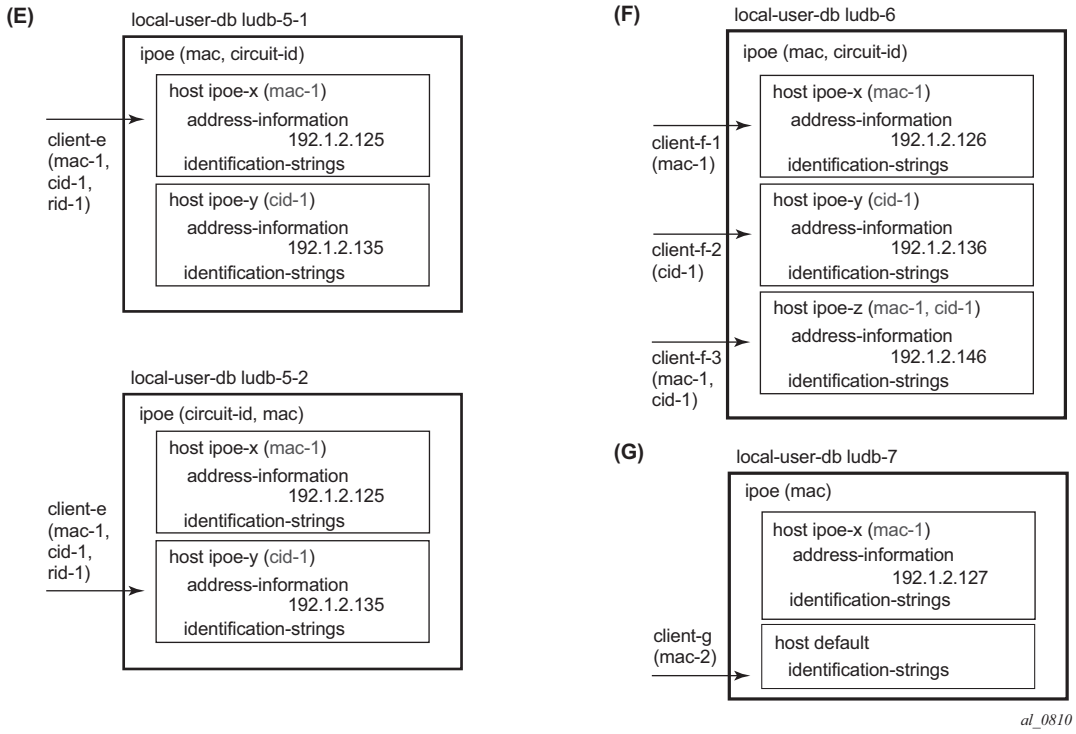

e. For the top part, matching is based on MAC-address and the circuit-id, in this sequence. When client-e connects (mac-1, cid-1, and rid-1), the system scans ludb-5-1 and matches host ipoe-x.

   For the bottom part, matching is based on the circuit-id first, then the MAC-address. When client-e connects (mac-1, cid-1, and rid-1), the system scans ludb-5-2 and matches host ipoe-y.

f. Matching is based on MAC-address and the circuit-id, in this sequence. When client-f-1 (mac-1) connects, the matching host is ipoe-x because only the MAC-address is provided and checked. When client-f-2 (cid-1) connects, the matching host is ipoe-y because only the client-id is provided and checked. When client-f-3 (mac-1, cid-1) connects, the matching host is ipoe-z.

g. Matching is based on the MAC-address only. When client-g with mac-2 connects, host default is matched because there is no explicit entry matching mac-2.

As shown is these examples, the system only checks the parameters provided by the client in the sequence as defined by the match-list. Parameters not provided by a client will not be searched for.

Tools Commands

The following **tools** command manually triggers the lookup of an IPoE host in an LUDB and is useful during commissioning, troubleshooting, and verification of the configured database, without the need for an external client.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-1" ipoe host-lookup
  - host-lookup [mac <ieee-address>] [remote-id <remote-id>|remote-id-hex <remote-
id-hex>] [sap-id <sap-id>] [service-id <service-id>] [string <vso-string>] [system-
id
    <system-id>] [option60 <option60-ascii> | option60-hex <option60-hex>] [circuit-
id <circuit-id-ascii>|circuit-id-hex <circuit-id-hex>]

 <ieee-address>       : xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx
 <remote-id-ascii>    : [255 chars max]
 <sap-id>             : [255 chars max]
 <service-id>         : [1..2148007978]|<svc-name:64 char max>
 <vso-string>         : [255 chars max]
 <system-id>          : [255 chars max]
 <option-60-ascii>    : [32 chars max]
 <circuit-id-ascii>   : [127 chars max]
 <circuit-id-hex>     : [0x0..0xFFFFFFFF...(max 254 hex nibbles)]
 <option60-hex>       : [0x0..0xFFFFFFFF...(max 64 hex nibbles)]
 <remote-id-hex>      : [0x0..0xFFFFFFFF...(max 510 hex nibbles)]
 <derived-id>         : [255 chars max]
```

A similar command exists for the lookup of a PPP host in an LUDB.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-1" ppp host-lookup
  - host-lookup [circuit-id <circuit-id>] [circuit-id-hex <circuit-id-
hex>] [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>]

 <circuit-id>         : [127 chars max]
 <circuit-id-hex>     : [0x0..0xFFFFFFFF...(max 254 hex nibbles)]
 <ieee-address>       : xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx
 <remote-id>          : [255 chars max]
 <remote-id-hex>      : [0x0..0xFFFFFFFF...(max 510 hex nibbles)]
 <sap-id>             : [255 chars max]
 <service-name>       : [255 chars max]
 <user-name>          : [253 chars max]
```

Example 1: Single Match Criterion

The following shows an excerpt from ludb-1. Host entry-11 defines the parameters for an IPoE host, and host entry-55 defines the parameters for a PPPoE host. Host matching IPoE hosts is MAC-address based, whereas host matching PPP hosts is username based.

```
configure
    subscriber-mgmt
        local-user-db "ludb-1" create
            description "example user-db"
```

```
            ipoe
                match-list mac
                host "default" create
                    address pool "pool4-1"
                    no shutdown
                exit
                host "entry-11" create
                    host-identification
                        mac 00:00:00:11:11:11
                    exit
                    address 10.1.1.211
                    --- snip ---
                    no shutdown
                exit
                --- snip ---
            exit
            ppp
                match-list username
                host "entry-55" create
                    host-identification
                        username "sub55@domain1"
                    exit
                    password chap sub55
                    address 10.1.2.252
                    --- snip ---
                    no shutdown
                exit
                --- snip ---
            exit
            no shutdown
        exit
```

IPoE hosts

IPoE host entry lookup using a MAC-address only is triggered with following **tools** command.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-1" ipoe host-
lookup mac 00:00:00:11:11:11
===============================================================================
IPoE Host Lookup results
===============================================================================
Result               : Success
Matched Host Name    : entry-11
Admin State          : Up
Last Mgmt Change     : 05/13/2015 05:19:50

Host Identification
 Circuit Id          : N/A
 Mac Address         : 00:00:00:11:11:11
 Remote Id           : N/A
 Sap Id              : N/A
 Service Id          : N/A
 String              : N/A
 Option 60           : N/A
 System Id           : N/A
 Encap Tag Range     : N/A
```

```
Matched Objects      : mac
 --- snipped ---
===============================================================================
*A:BNG-1#
```

The debug output confirms the successful lookup.

```
2546 2015/05/19 09:05:16.58 UTC MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  mac: 00:00:00:11:11:11

  Host entry-11 found in user data base ludb-1"
```

The following command is using a MAC-address, a circuit-id, and a remote-id for the lookup. The output shows that only the MAC-address is used, the other input parameters are ignored (N/A) so again entry-11 is selected.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-1" ipoe host-
lookup mac 00:00:00:11:11:11 circuit-id AA remote-id BB
===============================================================================
IPoE Host Lookup results
===============================================================================
Result               : Success
Matched Host Name    : entry-11
Admin State          : Up
Last Mgmt Change     : 05/13/2015 05:19:50

Host Identification
 Circuit Id          : N/A
 Mac Address         : 00:00:00:11:11:11
 Remote Id           : N/A
 Sap Id              : N/A
 Service Id          : N/A
 String              : N/A
 Option 60           : N/A
 System Id           : N/A
 Encap Tag Range     : N/A

Matched Objects      : mac
 --- snipped ---
===============================================================================
*A:BNG-1#
```

The following command triggers the lookup of a non-existing MAC-address, leading to a host not found message.

```
*A:BNG-1#  tools perform subscriber-mgmt local-user-db "ludb-1" ipoe host-
lookup mac 00:00:00:12:34:56
===============================================================================
IPoE Host Lookup results
===============================================================================
Result               : host not found
*A:BNG-1#
```

The host not found message is also confirmed by the debug output.

```
2550 2015/05/19 09:12:16.28 UTC MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host not found
  mac: 00:00:00:12:34:56

  Host not found in user data base ludb-1"
```

To allow IPoE users with unknown MAC-addresses to successfully connect, a default host can be created, at which time an informational message is returned:

```
*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe>host# host default create
                    address pool "pool4-1"
*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe>host# no shutdown
INFO: DHCPS #1138 This host will be considered as the default host
*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe>host#
```

After the previous commands are executed, devices with MAC-addresses not explicitly listed in the LUDB can also connect.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-1" ipoe host-
lookup mac 00:00:00:12:34:56
===============================================================================
IPoE Host Lookup results
===============================================================================
Result               : Success
Matched Host Name    : default
Admin State          : Up
Last Mgmt Change     : 05/19/2015 09:22:01

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
 Encap Tag Range     : N/A

Matched Objects      : N/A
--- snipped ---
===============================================================================
*A:BNG-1#
```

PPP hosts

Manually authenticating a PPP host is done as follows.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-
1" ppp authentication user-name sub55@domain1 password sub55
===============================================================================
Authentication results
===============================================================================
Result               : Success
Matched Host Name    : entry-55
```

```
Admin State          : Up
Last Mgmt Change     : 05/19/2015 09:32:42

Host Identification
 Mac Address         : N/A
 Circuit Id          : N/A
 Remote Id           : N/A
 Sap Id              : N/A
 Service Name        : N/A
 User Name           : sub55@domain1
 Encap Tag Range     : N/A

Matched Objects      : userName
--- snipped ---
===============================================================================
*A:BNG-1#
```

When the wrong password is provided, the following message is returned:

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-
1" ppp authentication user-name sub55@domain1 password sub5x
===============================================================================
Authentication results
===============================================================================
Result               : invalid password
*A:BNG-1#
```

PPP host entry lookup is similar to the IPoE host lookup. The following example demonstrates a user-name based lookup.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-1" ppp host-lookup user-
name sub55@domain1
===============================================================================
PPP host Lookup results
===============================================================================
Result               : Success
Matched Host Name    : entry-55
Admin State          : Up
Last Mgmt Change     : 05/19/2015 09:32:42

Host Identification
 Mac Address         : N/A
 Circuit Id          : N/A
 Remote Id           : N/A
 Sap Id              : N/A
 Service Name        : N/A
 User Name           : sub55@domain1
 Encap Tag Range     : N/A

Matched Objects      : userName
--- snipped ---
===============================================================================
*A:BNG-1#
```

The following command is using a user-name and a MAC-address for the lookup.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-1" ppp host-lookup user-
name sub55@domain1 mac 00:00:00:11:11:11
===============================================================================
PPP host Lookup results
===============================================================================
Result               : Success
Matched Host Name    : entry-55
Admin State          : Up
Last Mgmt Change     : 05/19/2015 09:32:42

Host Identification
 Mac Address         : N/A
 Circuit Id          : N/A
 Remote Id           : N/A
 Sap Id              : N/A
 Service Name        : N/A
 User Name           : sub55@domain1
 Encap Tag Range     : N/A

Matched Objects      : userName
--- snipped ---
===============================================================================
*A:BNG-1#
```

Similar to the IPoE host lookup, the lookup of a non-existing user fails if no default entry is defined for PPP. In this case, a default host can be defined.

Example 2: Multiple Match Criteria

The following shows an excerpt from ludb-2, with multiple match criteria.

The match-list includes mac, circuit-id, and remote-id, in this sequence.

```
configure
    subscriber-mgmt
        local-user-db "ludb-2" create
            ipoe
                match-list mac circuit-id remote-id
                host "entry-11" create
                    host-identification
                        mac 00:00:00:11:11:11
                    exit
                    address 10.1.1.111
                    --- snipped ---
                    no shutdown
                exit
                host "entry-12" create
                    host-identification
                        circuit-id string "11"
                        mac 00:00:00:11:11:11
                    exit
                    address 10.1.1.112
                    --- snipped ---
                    no shutdown
                exit
                host "entry-13" create
                    host-identification
```

```
                circuit-id string "11"
                mac 00:00:00:11:11:11
                remote-id string "AA"
            exit
            address 10.1.1.113
            --- snipped ---
            no shutdown
        exit
        host "entry-14" create
            host-identification
                circuit-id string "11"
                remote-id string "AA"
            exit
            address 10.1.1.114
            --- snipped ---
            no shutdown
        exit
```

The **tools** command below uses a MAC-address only, with entry-11 being matched.

```
*A:BNG-1#  tools perform subscriber-mgmt local-user-db "ludb-2" ipoe host-
lookup mac 00:00:00:11:11:11
===============================================================================
IPoE Host Lookup results
===============================================================================
Result               : Success
Matched Host Name    : entry-11
Admin State          : Up
Last Mgmt Change     : 05/19/2015 10:01:59

Host Identification
 Circuit Id          : N/A
 Mac Address         : 00:00:00:11:11:11
 Remote Id           : N/A
 Sap Id              : N/A
 Service Id          : N/A
 String              : N/A
 Option 60           : N/A
 System Id           : N/A
 Encap Tag Range     : N/A

Matched Objects      : mac
--- snipped ---
===============================================================================
*A:BNG-1#
```

The corresponding debug output shows the parameters from the match-list and their values, in sequence. The values for the circuit-id and the remote-id are left empty as they were not provided for the lookup.

```
2561 2015/05/19 10:22:29.55 UTC MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  mac: 00:00:00:11:11:11
  circuit-id:
  remote-id:

  Host entry-11 found in user data base ludb-2"
```

The following **tools** command uses a circuit-id and a remote-id, with entry-14 being matched.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-2" ipoe host-
lookup circuit-id 11 remote-id AA
===============================================================================
IPoE Host Lookup results
===============================================================================
Result               : Success
Matched Host Name    : entry-14
Admin State          : Up
Last Mgmt Change     : 05/19/2015 13:58:09

Host Identification
 Circuit Id          : 11
 Mac Address         : N/A
 Remote Id           : AA
 Sap Id              : N/A
 Service Id          : N/A
 String              : N/A
 Option 60           : N/A
 System Id           : N/A
 Encap Tag Range     : N/A

Matched Objects      : circ-id remote-id
 --- snipped ---
===============================================================================
*A:BNG-1#
```

The corresponding debug output shows that the original and the masked values of the circuit-id and the remote-id are the same, because no masks are applied.

```
2571 2015/05/19 14:10:00.18 UTC MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  mac:
  circuit-id:
    original:  11
    masked:    11
  remote-id:
    original:  AA
    masked:    AA

  Host entry-14 found in user data base ludb-2"
```

Accessing ludb-2 with a remote-id only returns a failure.

```
*A:BNG-1#  tools perform subscriber-mgmt local-user-db "ludb-2" ipoe host-
lookup remote-id AA
===============================================================================
IPoE Host Lookup results
===============================================================================
Result               : host not found
*A:BNG-1#

2564 2015/05/19 11:18:42.01 UTC MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host not found
```

```
  mac:
  circuit-id:
  remote-id:
    original:  AA
    masked:    AA

  Host not found in user data base ludb-2"
```

Example 3: Masking (1)

The following shows an excerpt from ludb-3, applying masks.

The match-list includes the circuit-id and the MAC-address, in this sequence. Masking applies to the circuit-id, which has the leading 8 characters and the trailing characters (behind the last vertical bar, and including the vertical bar) stripped.

```
configure
    subscriber-mgmt
        local-user-db "ludb-3" create
            description "masking example, ipoe"
            ipoe
                match-list circuit-id mac
                mask type circuit-id prefix-length 8 suffix-string "|*"
                host "entry-111" create
                    host-identification
                        mac 00:00:00:11:11:11
                        circuit-id string "grp-int-1-1"
                    exit
                    --- snipped ---
                    no shutdown
                exit
                --- snipped ---
```

The following **tools** command uses circuit-id and mac-address, matching entry-111.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-3" ipoe host-
lookup mac 00:00:00:11:11:11 circuit-id "BNG-1|1|grp-int-1-1|1/1/2/1:111"
===============================================================================
IPoE Host Lookup results
===============================================================================
Result               : Success
Matched Host Name    : entry-111
Admin State          : Up
Last Mgmt Change     : 05/20/2015 11:45:31

Host Identification
 Circuit Id          : grp-int-1-1
 Mac Address         : 00:00:00:11:11:11
 Remote Id           : N/A
 Sap Id              : N/A
 Service Id          : N/A
 String              : N/A
 Option 60           : N/A
 System Id           : N/A
 Encap Tag Range     : N/A
```

```
Matched Objects      : circ-id mac
--- snipped ---
===============================================================================
*A:BNG-1#
```

The debug output shows the values of the parameters before and after applying the mask.

```
2576 2015/05/20 11:50:25.13 UTC MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  circuit-id:
    original:  BNG-1|1|grp-int-1-1|1/1/2/1:111
    masked:    grp-int-1-1
  mac: 00:00:00:11:11:11

  Host entry-111 found in user data base ludb-3"
```

Example 4: Masking (2)

The following shows an excerpt from ludb-4, applying masks.

The match-list includes the username, circuit-id, and remote-id, in this sequence. Masking applies to both the username and the circuit-id. The username has everything before the **@**-sign and the trailing .org stripped. The circuit-id has the trailing 11 characters stripped.

```
configure
    subscriber-mgmt
        local-user-db "ludb-4" create
            ppp
                match-list username circuit-id remote-id
                mask type circuit-id suffix-length 11
                mask type username prefix-string "*@" suffix-string ".org"
                host "entry-11" create
                    host-identification
                         username domain1
                         circuit-id string "BSAN-2"
                    exit
                    address pool "pool4-1"
                    identification-strings 254 create
                        sla-profile-string "sla-prof-1"
                        sub-profile-string "sub-prof-2"
                    exit
                    no shutdown
                exit
                --- snipped ---
            exit
            no shutdown
        exit
        --- snipped ---
```

The following **tools** command does not result in a match, which is not the intention.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-4" ppp host-lookup user-
```

```
name sub11@domain1.org circuit-id "BSAN-2|100|1/2/1:111"
===============================================================================
PPP host Lookup results
===============================================================================
Result               : host not found
*A:BNG-1#
```

The debug output shows the original and the masked value of the user-name and the circuit-id; the remote-id was not provided.

```
2581 2015/05/20 12:12:57.00 UTC MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host not found
  user-name:
    original:  sub11@domain1.org
    masked:    domain1
  circuit-id:
    original:  BSAN-2|100|1/2/1:111
    masked:    BSAN-2|10
  remote-id:

  Host not found in user data base ludb-4"
```

The preceding output shows that three more characters (|10) must be stripped to have a successful lookup, and following configuration changes are needed.

```
configure
    subscriber-mgmt
        local-user-db "ludb-4" create
            ppp
                mask type circuit-id suffix-length 14
```

Modifying the mask results in host entry-11 being matched.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-4" ppp host-lookup user-
name sub11@domain1.org circuit-id "BSAN-2|100|1/2/1:111"
===============================================================================
PPP host Lookup results
===============================================================================
Result               : Success
Matched Host Name    : entry-11
Admin State          : Up
Last Mgmt Change     : 05/20/2015 12:03:42

Host Identification
 Mac Address         : N/A
 Circuit Id          : BSAN-2
 Remote Id           : N/A
 Sap Id              : N/A
 Service Name        : N/A
 User Name           : domain1
 Encap Tag Range     : N/A

Matched Objects      : userName circ-id
--- snipped ---
===============================================================================
*A:BNG-1#
```

The debug output shows the effect of the modified mask.

```
231 2015/05/29 07:09:33.05 UTC MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  user-name:
    original:  sub11@domain1.org
    masked:    domain1
  circuit-id:
    original:  BSAN-2|100|1/2/1:111
    masked:    BSAN-2
  remote-id:

  Host entry-11 found in user data base ludb-4"
```

Example 5: VLAN Range

LUDB matching also supports the use of encap-tag-range. Host-identification then needs a start and an end for the range, which both use the following format:

```
dot1q          - qtag1
    qinq           - (qtag1.qtag2 | qtag1.* | *.qtag2)
    atm            - (vpi/vci | vpi/* | */vci)
      qtag1          - [0..4094]
      qtag2          - [0..4094]
      vpi            - [0..4095] (NNI)
                       [0..255] (UNI)
      vci            - [1..65535]
```

For VLAN tagging, the Ethernet frames could be single- or dual- tagged. For ATM, a virtual path identifier (VPI) and a virtual circuit identifier (VCI) can be defined. The asterisk (*) serves as a wildcard, meaning that the parameter is ignored.

The system validates, at configuration time, the values of the start-tag and the end-tag, applying following rules:

- The start-tag must be lower than the end-tag.
- When using the asterisk, it should be present in both the start-tag and the end-tag:
  - *.10 - *.100 — the outer tag is ignored
  - 201.* - 299.* — the inner tag is ignored
- The encapsulation type for start-tag and end-tag must be the same.
- Overlapping ranges (while on the same port) are not allowed.

The following shows an excerpt from ludb-5, using vlan-ranges.

```
configure
    subscriber-mgmt
        local-user-db "ludb-5" create
            description "example for vlan ranges"
            ipoe
```

```
                match-list encap-tag-range
                host "range-1" create
                    host-identification
                        encap-tag-range start-tag *.1 end-tag *.50
                    exit
                    address pool "pool4-3"
                    --- snipped ---
                    no shutdown
                exit
                host "range-2" create
                    host-identification
                        encap-tag-range start-tag *.51 end-tag *.100
                    exit
                    address pool "pool4-4"
                    --- snipped ---
                    no shutdown
                exit
            exit
            no shutdown
        exit
```

The following **tools** command specifies a sap-id including an outer and an inner tag, matching host range-1.

```
*A:BNG-1# tools perform subscriber-mgmt local-user-db "ludb-5" ipoe host-lookup sap-
id 1/1/1:50.4
===============================================================================
IPoE Host Lookup results
===============================================================================
Result               : Success
Matched Host Name    : range-1
Admin State          : Up
Last Mgmt Change     : 05/20/2015 14:03:10

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
 Encap Tag Range     : start-tag *.1 end-tag *.50

Matched Objects      : encap-tag-range

Address              : pool "pool4-3"
--- snipped ---
===============================================================================
*A:BNG-1#
```

Operational Considerations

- Names of LUDBs and LUDB entries cannot be changed.
- Modification of the host identification fields is possible only when the host-entry is put in the shutdown state. Modifying output fields does not require the host-entry to be in the shutdown state.

```
*A:BNG-1>config>subscr-mgmt>loc-user-db>ppp>host>host-ident# circuit-id string
            x-y-z
MINOR: DHCPS #1133 Not allowed. Host is not shutdown
```

- Modifying a match-list requires the LUDB to be in the shutdown state.
- Modifying a match-list results in a re-evaluation of all host entries of the LUDB block, so that the lookup database and the unmatched host list are re-populated.

```
*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe# match-list circuit-id
INFO: DHCPS #1107 Host could not be inserted in lookup database - lookup database
            constructed, 1 hosts not inserted: 1 no match, 0 duplicate
```

  Modifying the match-list also imposes the risk of a default entry with host-identification fields suddenly not being the fallback (default) entry anymore, which is why defining a default entry with host-identification fields is not recommended.
- Modifying one or more mask types does not require the LUDB to be in the shutdown state.
- Deletion of an LUDB requires that the LUDB is not referenced and the LUDB is in the shutdown state. Use caution: the status of the individual entries is not taken into account when deleting an LUDB.

```
*A:BNG-1>config>subscr-mgmt# no local-user-db "ludb-1"
MINOR: DHCPS #1103 User data base still referenced

*A:BNG-1>config>subscr-mgmt# no local-user-db "ludb-11"
MINOR: DHCPS #1104 Not allowed when user db admin state is up
```

Troubleshooting and Debugging LUDBs

The **tools** command can also be used for troubleshooting; the example is not repeated for brevity.

Show Commands

The following command shows which LUDBs are available in the system, including the administrative state and the host count. The host count equals the total number of configured ipoe and ppp entries, regardless of their administrative state (shutdown/no shutdown).

```
*A:BNG-1# show subscriber-mgmt local-user-db
===============================================================================
Local User Databases
===============================================================================
Name                            Admin Host   Description
                                State Count
-------------------------------------------------------------------------------
ludb-1                          Up    9      users db for sub-int-1
ludb-2                          Down  1
ludb-3                          Down  6      test for pppoe users
-------------------------------------------------------------------------------
Number of Local User Databases : 3     Number of Hosts : 16
===============================================================================
*A:BNG-1#
```

For showing the host count and the IPoE and PPP match types for a single LUDB, following command is useful.

```
*A:BNG-1# show subscriber-mgmt local-user-db "ludb-1"
===============================================================================
Local User Database "ludb-1"
===============================================================================
Description             : users db for sub-int-1
Admin State             : Up
Last Mgmt Change        : 03/12/2015 15:54:46
Host Count              : 9
IPoE Match Types        : mac
PPP Match Types         : userName
===============================================================================
*A:BNG-1#
```

Listing all IPoE hosts in a specific LUDB is performed with the following command.

```
*A:BNG-1# show subscriber-mgmt local-user-db "ludb-1" ipoe-all-hosts
===============================================================================
Local User Database "ludb-1" IPoE hosts
===============================================================================
Name                            Admin     Matched objects
                                State
-------------------------------------------------------------------------------
mac-11                          Up        mac
--- snipped ---
default                         Up        -
-------------------------------------------------------------------------------
Number of IPoE Hosts : 5
===============================================================================
*A:BNG-1#
```

A similar command lists all PPP hosts.

```
*A:BNG-1# show subscriber-mgmt local-user-db "ludb-1" ppp-all-hosts
===============================================================================
Local User Database "ludb-1" PPP Hosts
===============================================================================
Name                            Admin     Matched objects
                                State
-------------------------------------------------------------------------------
ppp-55                          Up        userName
--- snipped ---
-------------------------------------------------------------------------------
Number of PPP Hosts : 4
===============================================================================
*A:BNG-1#
```

To find the places where a specific LUDB is applied, use the following command.

```
*A:BNG-1# show subscriber-mgmt local-user-db "ludb-1" association
No DHCP Server associations found.
===============================================================================
DHCP client interface associations for ludb-1
===============================================================================
Interface-Name                          Svc-Id     Type
-------------------------------------------------------------------------------
grp-int-1-1                             1          IES
grp-int-1-2                             1          IES
grp-int-2-1                             1          IES
grp-int-2-2                             1          IES
-------------------------------------------------------------------------------
No. of Interface(s): 4
===============================================================================

===============================================================================
DHCP6 interface associations for ludb-1
===============================================================================
Interface-Name                          Svc-Id     Type
-------------------------------------------------------------------------------
grp-int-1-1                             1          IES
grp-int-1-2                             1          IES
grp-int-2-1                             1          IES
grp-int-2-2                             1          IES
-------------------------------------------------------------------------------
No. of Interface(s): 4
===============================================================================

No Router solicit interface associations found.
No PPP client interface associations found.

===============================================================================
PPPoE client interface associations for ludb-1
===============================================================================
Interface-Name                          Svc-Id     Type
-------------------------------------------------------------------------------
grp-int-1-1                             1          IES
grp-int-1-2                             1          IES
grp-int-2-1                             1          IES
grp-int-2-2                             1          IES
```

```
-------------------------------------------------------------------------------
No. of Interface(s): 4
===============================================================================

No IPoE client interface associations found.

No capture SAP associations found.

No associated L2TP groups found.

No associated L2TP tunnels found.

No associated authentication policies found.

No WPP interface associations found.
*A:BNG-1#
```

The following command is useful for displaying the details of a specific LUDB entry.

```
*A:BNG-1# show subscriber-mgmt local-user-db "ludb-1" ipoe-host "mac-33"
===============================================================================
IPoE Host "mac-33"
===============================================================================
Admin State          : Up
Last Mgmt Change     : 03/12/2015 15:54:46

Host Identification
 Circuit Id          : N/A
 Mac Address         : 00:00:00:33:33:33
 Remote Id           : N/A
 Sap Id              : N/A
 --- snipped ---
 Matched Objects      : mac

Address              : 10.1.2.233
Auth Policy          : N/A
Acct Policy          : N/A
Dupl Acct Policy     : N/A
Auth Domain Name     : N/A
Diameter app policy  : (Not Specified)
Diameter auth policy : (Not Specified)
Rip Policy           : N/A
IPv6 Address         : 2001:db8:102:33::3333
IPv6 Del Pfx         : 2001:db8:f102:3300::/56
--- snipped ---
IPv6 Del Pfx Length  : 56

--- snipped ---

Identification Strings (option 254)
 Subscriber Id       : ludb-1-33
 SLA Profile String  : sla-prof-2

--- snipped ---

Filter Overrules
 Ing Ipv4 Fltr       : N/A
 Egr Ipv4 Fltr       : N/A
```

```
 Ing Ipv6 Fltr       : N/A
 Egr Ipv6 Fltr       : N/A
===============================================================================
*A:BNG-1#
```

The following commands list the IPoE and the PPP host entries in a specific LUDB that are not matched. Note that duplicates are counted as unmatched hosts.

```
*A:BNG-1# show subscriber-mgmt local-user-db "ludb-22" ipoe-unmatched-hosts
===============================================================================
Local User Database "ludb-22" IPoE unmatched hosts
===============================================================================
Name                              Reason      Duplicate Host
-------------------------------------------------------------------------------
this-is-a-no-match                No match    N/A
this-is-a-duplicate               Duplicate   entry-12
-------------------------------------------------------------------------------
Number of IPoE Unmatched Hosts : 2
===============================================================================
*A:BNG-1#


*A:BNG-1# show subscriber-mgmt local-user-db "ludb-22" ppp-unmatched-hosts
===============================================================================
Local User Database "ludb-22" PPP unmatched hosts
===============================================================================
Name                              Reason      Duplicate Host
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No PPP Unmatched Hosts found
===============================================================================
*A:BNG-1#
```

Debugging Commands

The following configuration enables debugging for both ludb-1 and for ludb-2.

```
debug
    subscriber-mgmt
        local-user-db "ludb-1"
            detail all
        exit
        local-user-db "ludb-2"
            detail all
        exit
    exit
exit
```

To ensure that the debug output is sent to the console, the following additional configuration is needed.

```
configure
    log
```

```
        log-id 1
            from debug-trace
            to session
            no shutdown
        exit
    exit
```

After the preceding configuration, debug output appears as part of the session.

```
314 2015/03/16 14:56:40.52 UTC MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  mac: 00:00:00:aa:aa:aa

  Host default found in user data base ludb-1"
```

Conclusion

In this example general LUDB concepts are explained. LUDBs are defined and host entries for both IPoE as well as for PPP are described. The different match criteria are explained and demonstrated by means of examples, including the use of single and multiple match criteria. Match criteria are handled left to right, in sequence, so that a natural priority is taken care of. Debugging aids are provided through **show**, **debug** and **tools** commands.

Local User Database for DHCPv4 Server

This chapter provides information about Local User Database (LUDB) for DHCPv4 server.

Topics in this chapter include:

Applicability

This chapter is applicable to the 7x50 SR series and was tested on SR OS 13.0.R1. Chassis mode B or later must be used.

Basic LUDB knowledge is a prerequisite for understanding this chapter.

Overview

In SR OS, a local DHCPv4 server can be assigned a Local User Database (LUDB).

Assigning an LUDB to a DHCPv4 server allows the server to:

- control IP address assignment; for example, by assigning a fixed IP address based on the user's MAC address
- control DHCPv4 options for native as well as for simulated DHCPv4 clients used by PPP. In the case of PPP, users are identified to the DHCPv4 server using the DHCPv4 Vendor-Specific Information Sub-option [82,9][6] in the DHCPv4 discover/request messages
- provide ESM strings (referred to as identification-strings in CLI) using a user-defined unassigned DHCPv4 option (option 254 is provided as default)

Introduction

A local DHCPv4 server can be used for IPoE users as well as for PPP users (see Figure 188).

When a DHCPv4 user connects, the typical Discover, Offer, Request, and Acknowledge (DORA) message sequence running between the DHCPv4 client and the DHCPv4 server also passes through a relay agent.

When a PPP user connects through LCP/IPCP, an internal DHCPv4 client manages the communication toward the DHCPv4 server, on the condition that the relay agent also has relaying enabled for PPP applications. This internal DHCPv4 client is also referred to as a local (DHCPv4) client.

The DHCPv4 server can be located in the same node as the relay agent and the local client, but that is not required.

Figure 188 **LUDB Access via a DHCPv4 Server**

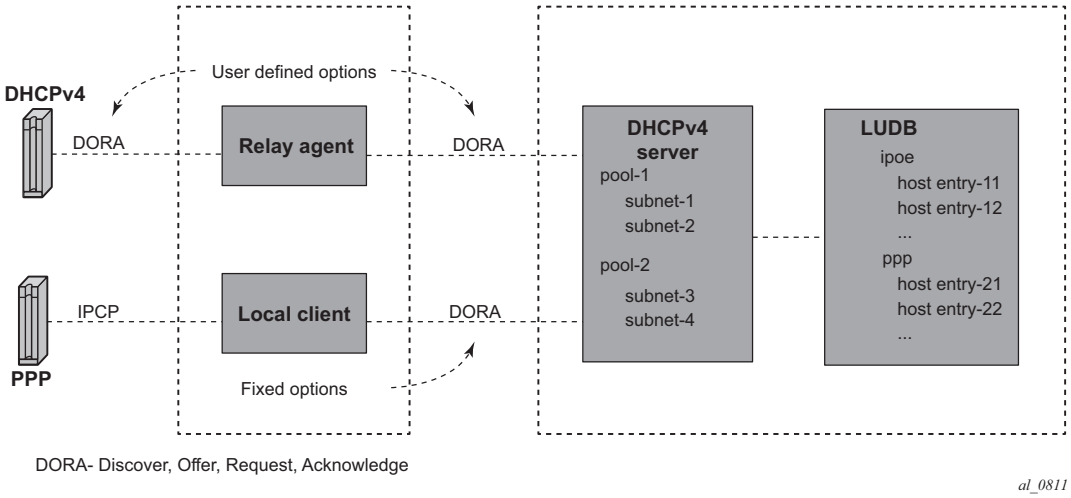

The relay agent must be configured correctly in order to forward the messages toward a DHCPv4 server. The DHCPv4 server IP address is defined. The relay agent can include following configurable options and sub-options to be used by the DHCPv4 server:

- [82,1] Agent Circuit ID Sub-option
- [82,2] Agent Remote ID Sub-option
- [82,9] Vendor-Specific Information Sub-option (VSO)
  - [1] system-ID (7x50 system-name)
  - [2] client MAC-address

  - [3] service-ID (7x50 IES/VPRN service-id)
  - [4] sap-ID
  - [5] string
  - [13] pool-name

The local client uses the same DHCPv4 server IP address as the preceding relay agent, and can include its own set of unconfigurable options and sub-options to be used by the DHCPv4 server:

- [60] vendor class (fixed string: ALU7XXXSBM)
- [82,1] Agent Circuit ID Sub-option
- [82,2] Agent Remote ID Sub-option
- [82,6] Subscriber-ID Sub-option (equals PPPoE username)
- [82,9] Vendor-Specific Information Sub-option
  - [1] system-ID (not included in a redundant node)
  - [2] client MAC-address
  - [3] service-ID (7x50 IES/VPRN service-id)
  - [4] sap-ID (not included for retail VPRN and redundant node)
  - [6] client type (1=ppp)
  - [13] pool-name
  - [14] service name (PPPoE tag service-name)
  - [17] session-ID (PPPoE session-id)

When one or more of these options and sub-options are included, the DHCPv4 server can use them while accessing the LUDB, for selection of the section (client type 1 is the PPP section), and the host entry in that section. For example, not including the service-ID VSO [82,9][3] to the DHCPv4 server, when LUDB host identification needs the MAC address and service ID, will result in an LUDB lookup failure on the DHCPv4 server and a silent drop of the DHCPv4 discover message.

LUDB Input Parameters

The following IPoE host identification fields are supported when accessing an LUDB from a DHCPv4 server:

- circuit-id
- encap-tag-range
- mac

- option60
- remote-id
- sap-id
- service-id
- string
- system-id

The LUDB lookup process can match up to four IPoE match-criteria, as defined by the IPoE match-list.

The following PPP host identification fields are supported when accessing an LUDB from a DHCPv4 server:

- circuit-id
- encap-tag-range
- mac
- remote-id
- sap-id
- service-name
- username

The LUDB lookup process can match up to three PPP match-criteria, as defined by the PPP match-list.

LUDB Output Parameters

Addressing Information

The host entry address field has the following configuration options when the LUDB is associated with a DHCPv4 server:

- no address

  Host access is not allowed. The clients mapping to this host entry will not get an IP address.

```
3874 2015/05/27 11:44:52.84 UTC MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4
DISCOVER dropped: host=entry-55, host found but no valid address info
```

- address <*ip-address*>

  A fixed IP address is offered to the client and should not overlap with the address ranges configured in the local DHCP server.

- address pool <*pool-name*> [secondary-pool <*sec-pool-name*>]

  The DHCPv4 server allocates an address from one of the subnets in that pool on the condition that the DHCPv4 server is configured to use pool names for address selection. Optionally, a secondary pool can be defined, to be used in case the primary pool is exhausted.

  Pool-name addressing is useful when the subscriber management node is not capable of inserting the pool-name VSO [82,9][13] or when a specific host requires a specific pool-name different from the pool-name included by the subscriber management node.

- address gi-address [scope <*subnet*|*pool*>]

  When the scope is set to subnet, the DHCPv4 server allocates an address from the subnet that includes the Gi address. When the scope is set to pool, the DHCPv4 server allocates an address from the subnet that includes the Gi address, or from the other subnets belonging to the same pool. Gi-addressing is useful when the subnet that the Gi address belongs to is exhausted.

- address use-pool-from-client [delimiter <*delimiter*>]

  The DHCPv4 server allocates an address from one of the subnets in the pool, as indicated by the pool-name VSO [82,9][13]. If two pools are available in this VSO, the configured delimiter distinguishes the first pool-name from the second pool-name.

Identification Strings

An LUDB can optionally return identification strings (also known as ESM strings). The DHCPv4 server returns them in a user-defined DHCPv4 option (default: 254) to the requesting entity (the relay agent or the local client). The identification strings, in alphabetical order, are:

- ancp-string
- app-profile-string
- category-map-name
- inter-dest-id
- sla-profile-string
- sub-profile-string
- subscriber-id

Options

An LUDB can return options to be used by the relay agent, the internal client, or the end-user device.

Available IPoE user options, configurable by option-name are:

- default-router
- dns-server
- domain-name
- lease-rebind-timer
- lease-renew-timer
- lease-time
- netbios-name-server
- netbios-node-type
- subnet-mask

Available PPP user options, configurable by option-name are:

- dns-server
- netbios-name-server

Additional options, configurable by option-number, can be configured for both IPoE and PPP users by using the **custom-option** command.

```
*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe>host>options# custom-option
  - custom-option <option-number> address [<ip-address>...(upto 4 max)]
  - custom-option <option-number> hex <hex-string>
  - custom-option <option-number> string <ascii-string>
  - no custom-option <option-number>

 <option-number>      : [1..254]
 <ip-address>         : a.b.c.d
 <ascii-string>       : [127 chars max]
 <hex-string>         : [0x0..0xFFFFFFFF...(max 254 hex nibbles)]

*A:BNG-1>config>subscr-mgmt>loc-user-db>ipoe>host>options#
```

The encoding of these custom options is either in hexadecimal, ASCII, or IP address format. In debug output, these custom options are indicated as **Unknown options** and presented in Type-Length-Value (TLV) format.

Options and custom options can be configured at three different levels:

- LUDB host entry level
- DHCPv4 server pool level

- DHCPv4 server subnet level

Options and custom options defined at the host entry level overrule options defined at either of the server levels.

Other Parameters

All other parameters in the host entry definition are silently ignored (for example, IPv6 related parameters, msap-defaults, retail-service-id, and so on) because they are not applicable for a DHCPv4 server-associated LUDB.

Configuration

Figure 189 shows the example configuration used in this chapter.

Figure 189 **Example Configuration**

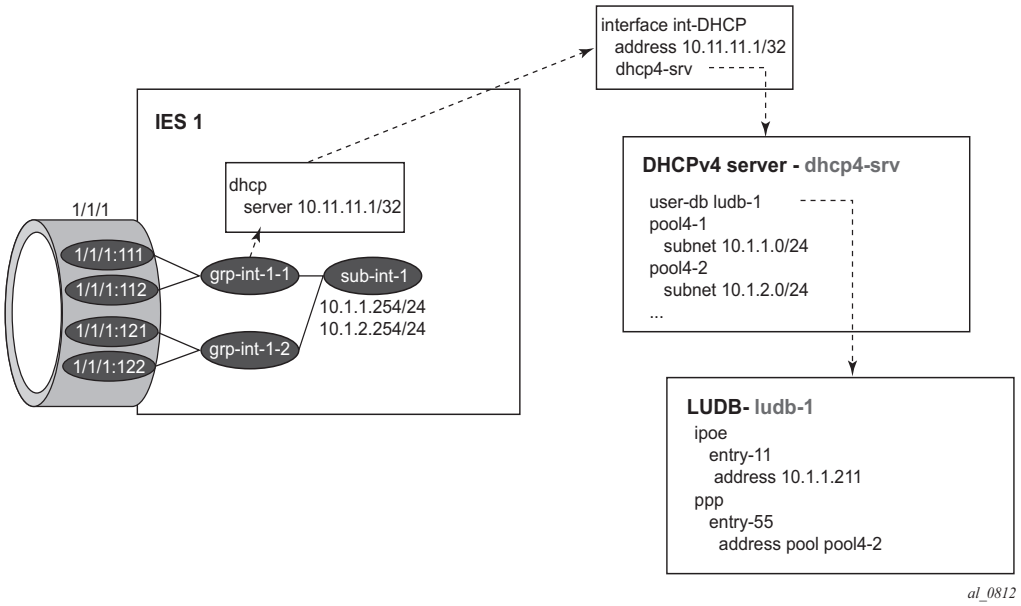

An LUDB can be associated with a DHCPv4 server in the base router instance or in a VPRN service instance using the following commands:

```
configure router dhcp local-dhcp-server <name> user-db <name>
configure service vprn <service-id> dhcp local-dhcp-server <name> user-db <name>
```

The following example has the DHCPv4 server created in the base router instance.

```
configure
    router
        dhcp
            local-dhcp-server "dhcp4-srv" create
                user-db "ludb-1"
                use-gi-address
                use-pool-from-client
                pool "pool4-1" create
                    subnet 10.1.1.0/24 create
                        options
                            subnet-mask 255.255.255.0
                            default-router 10.1.1.254
                        exit
                        address-range 10.1.1.1 10.1.1.100
                    exit
                exit
                pool "pool4-2" create
                    subnet 10.1.2.0/24 create
                        options
                            subnet-mask 255.255.255.0
                            default-router 10.1.2.254
                        exit

                        address-range 10.1.2.1 10.1.2.254
                    exit
                exit
                --- snipped ---
                no shutdown
            exit
        exit
```

The server is then associated with the loopback interface at address 10.11.11.1.

```
configure
    router
        interface "int-DHCP"
            address 10.11.11.1/32
            loopback
            local-dhcp-server "dhcp4-srv"
            no shutdown
        exit
    exit
```

The following is a partial configuration of service IES 1. Note that the DHCPv4 relay agent is configured to include the service-ID VSO [82,9][3] because this option is used in the LUDB for matching purposes.

```
configure
   service
        ies 1 customer 1 create
            subscriber-interface "sub-int-1" create
                address 10.1.1.254/24
                address 10.1.2.254/24
                group-interface "grp-int-1-1" create
```

```
                    arp-populate
                    dhcp
                        option
                            action replace
                            circuit-id
                            remote-id
                            vendor-specific-option
                                service-id
                            exit
                        exit
                        server 10.11.11.1
                        trusted
                        lease-populate 100
                        client-applications dhcp ppp
                        gi-address 10.1.1.254
                        no shutdown
                    exit
                    sap 1/1/1:111 create
                        sub-sla-mgmt
                            def-sub-profile "sub-prof-1"
                            def-sla-profile "sla-prof-1"
                            sub-ident-policy "sub-id-pol-1"
                            multi-sub-sap
                            no shutdown
                        exit
                    exit
                    --- snipped ---
                    pppoe
                        session-limit 100
                        sap-session-limit 100
                        no shutdown
                    exit
                exit
            exit
```

In the following example, IPoE users are matched against the service-id, the MAC address and option 60. Host entry-11 returns a fixed IP address, three identification strings, and a set of options. PPP users are matched against the MAC address. Host entry-55 returns an address-pool, two identification strings, and two DNS servers as options.

```
configure
    subscriber-mgmt
        local-user-db "ludb-1" create
            description "example user-db"
            ipoe
                match-list service-id mac option60
                host "entry-11" create
                    host-identification
                        mac 00:00:00:11:11:11
                        service-id 1
                        option60 hex 0xaabb
                    exit
                    address 10.1.1.211
                    identification-strings 254 create
                        subscriber-id "sub-11"
                        sla-profile-string "sla-profile-1"
```

```
                        sub-profile-string "sub-profile-1"
                    exit
                    options
                        subnet-mask 255.255.255.0
                        default-router 10.1.1.251
                        dns-server 2.2.2.2 2.2.2.1
                        domain-name "domain.org"
                        netbios-name-server 10.1.1.252
                        netbios-node-type B
                        lease-time hrs 12
                        custom-option 251 hex 0x010203
                    exit
                    no shutdown
                exit
                --- snip ---
            exit
            ppp
                match-list mac
                host "entry-55" create
                    host-identification
                        mac 00:00:00:55:55:55
                    exit
                    address pool "pool4-2"
                    identification-strings 254 create
                        subscriber-id "sub-55"
                        sla-profile-string "sla-prof-3"
                        sub-profile-string "sub-prof-3"
                    exit
                    options
                        dns-server 2.2.2.2 2.2.2.1
                    exit
                    --- snip ---
                    no shutdown
                exit
                --- snip ---
            exit
            no shutdown
        exit
```

Note that entry-11 defines a fixed IP address in one of the subnets allowed on the group interface grp-int-1-1 on IES 1, but out of the range defined in the DHCPv4 server.

Debugging

The following example debugs the local DHCP server and the LUDB ludb-1.

```
debug
    router "Base"
        local-dhcp-server "dhcp4"
            detail-level medium
            mode egr-ingr-and-dropped
        exit
```

```
    exit
    subscriber-mgmt
        local-user-db "ludb-1"
            detail all
        exit
    exit
exit
```

The following additional configuration ensures that the debug output is sent to the current login session.

```
configure
    log
        log-id 1
            from debug-trace
            to session
            no shutdown
        exit
    exit
exit
```

IPoE Users Verification

The following command shows the DHCPv4 server lease state record for LUDB host entry-11. The address type is set to **fixed** because ludb-1 returns a fixed IP address.

```
*A:BNG-1# show router dhcp local-dhcp-server "dhcp4" leases 10.1.1.211 detail
===============================================================================
Lease for DHCP server dhcp4 router Base
===============================================================================
IP-address                    : 10.1.1.211
Lease-state                   : stable
Lease started                 : 2015/06/05 14:16:34
Last renew                    : N/A
Remaining LifeTime            : 11h57m9s
Remaining Potential Exp. Time: 0h0m0s
MAC address                   : 00:00:00:11:11:11
Xid                           : 0x1
Failover Control              : local
Client Type                   : dhcp
User-db Host Name             : entry-11
User-db Address Type          : fixed
Persistence Key               : N/A
Opt82 Hex Dump                : (length=80)
                              : 52 4e 01 1d 42 4e 47 2d 31 7c 31 7c 67 72 70 2d
                              : 69 6e 74 2d 31 2d 31 7c 31 2f 31 2f 31 3a 31 31
                              : 31 02 06 00 00 00 11 11 11 09 25 00 00 19 7f 20
                              : 01 05 42 4e 47 2d 31 02 06 00 00 00 11 11 11 03
                              : 04 00 00 00 01 04 09 31 2f 31 2f 31 3a 31 31 31
Opt82 Circuit Id              : BNG-1|1|grp-int-1-1|1/1/1:111
Opt82 Remote Id               : (hex) 00 00 00 11 11 11
Opt82 Subscr Id               :
Opt82 VS System               : BNG-1
```

```
Opt82 VS Clnt MAC            : 00:00:00:11:11:11
Opt82 VS Service             : (hex) 00 00 00 01
Opt82 VS SAP                 : 1/1/1:111
Opt82 VS String              :
Opt60 Hex Dump               : (length=2)
                             : aa bb
Lease Remaining Hold Time    : 0h0m0s
===============================================================================
*A:BNG-1#
```

The debug output on the DHCPv4 server and the LUDB ludb-1 shows that the LUDB is accessed for every incoming message. Note that the identification-strings are returned to the relay agent in the Offer and Acknowledge messages through option [254]; see Figure 190 for the decoding.

```
94 2015/06/05 14:16:34.34 UTC MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4
Rx DHCP Discover

   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 10.1.1.254
   chaddr: 00:00:00:11:11:11    xid: 0x1

   DHCP options:
   [82] Relay agent information: len = 78
      [1] Circuit-id: BNG-1|1|grp-int-1-1|1/1/1:111
      [2] Remote-id: (hex) 00 00 00 11 11 11
      [9] Vendor-Specific info: len = 37
          Enterprise [6527] : len = 32
          [1] systemId: BNG-1
          [2] clntMac: 00:00:00:11:11:11
          [3] servId: 1
          [4] sapId: 1/1/1:111
   [53] Message type: Discover
   [60] Class id: (hex) aa bb
   [255] End
"

95 2015/06/05 14:16:34.34 UTC MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  svc-id: 1
  mac: 00:00:00:11:11:11
  option60:
    original:  0xaabb
    masked:    0xaabb
  encap-tag: 111

  Host entry-11 found in user data base ludb-1"

96 2015/06/05 14:16:34.34 UTC MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4
lease added for 10.1.1.211 state=offer
"

97 2015/06/05 14:16:34.35 UTC MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4
Tx DHCP Offer to local relay agent 10.1.1.254 vrId=1
```

```
   ciaddr: 0.0.0.0           yiaddr: 10.1.1.211
   siaddr: 10.11.11.1        giaddr: 10.1.1.254
   chaddr: 00:00:00:11:11:11    xid: 0x1

   DHCP options:
   [82] Relay agent information: len = 78
      [1] Circuit-id: BNG-1|1|grp-int-1-1|1/1/1:111
      [2] Remote-id: (hex) 00 00 00 11 11 11
      [9] Vendor-Specific info: len = 37
          Enterprise [6527] : len = 32
          [1] systemId: BNG-1
          [2] clntMac: 00:00:00:11:11:11
          [3] servId: 1
          [4] sapId: 1/1/1:111
   [53] Message type: Offer
   [54] DHCP server addr: 10.11.11.1
   [51] Lease time: 43200
   [1] Subnet mask: 255.255.255.0
   [3] Router: 10.1.1.251
   [6] Domain name server: length = 8
            2.2.2.2
            2.2.2.1
   [15] Domain name: domain.org
   [44] NETBIOS name server: 10.1.1.252
   [46] NETBIOS type: 1
   [251] Unknown option: len = 3, value = 01 02 03
   [254] Unknown option: len = 38, value = 07 06 73 75 62 2d 31 31 08 0d 73
   6c 61 2d 70 72 6f 66 69 6c 65 2d 31 09 0d 73 75 62 2d 70 72 6f 66 69 6c 65
   2d 31
   [60] Class id: (hex) aa bb
   [255] End
"

98 2015/06/05 14:16:34.44 UTC MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4
Rx DHCP Request

   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 10.1.1.254
   chaddr: 00:00:00:11:11:11    xid: 0x1

   DHCP options:
   [82] Relay agent information: len = 78
      [1] Circuit-id: BNG-1|1|grp-int-1-1|1/1/1:111
      [2] Remote-id: (hex) 00 00 00 11 11 11
      [9] Vendor-Specific info: len = 37
          Enterprise [6527] : len = 32
          [1] systemId: BNG-1
          [2] clntMac: 00:00:00:11:11:11
          [3] servId: 1
          [4] sapId: 1/1/1:111
   [53] Message type: Request
   [50] Requested IP addr: 10.1.1.211
   [60] Class id: (hex) aa bb
   [54] DHCP server addr: 10.11.11.1
   [255] End
"
```

```
99 2015/06/05 14:16:34.44 UTC MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  svc-id: 1
  mac: 00:00:00:11:11:11
  option60:
    original:  0xaabb
    masked:    0xaabb
  encap-tag: 111

  Host entry-11 found in user data base ludb-1"

100 2015/06/05 14:16:34.44 UTC MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4
lease update for 10.1.1.211 state=stable
"

101 2015/06/05 14:16:34.44 UTC MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4
Tx DHCP Ack to local relay agent 10.1.1.254 vrId=1

   ciaddr: 0.0.0.0           yiaddr: 10.1.1.211
   siaddr: 10.11.11.1        giaddr: 10.1.1.254
   chaddr: 00:00:00:11:11:11    xid: 0x1

   DHCP options:
   [82] Relay agent information: len = 78
      [1] Circuit-id: BNG-1|1|grp-int-1-1|1/1/1:111
      [2] Remote-id: (hex) 00 00 00 11 11 11
      [9] Vendor-Specific info: len = 37
          Enterprise [6527] : len = 32
          [1] systemId: BNG-1
          [2] clntMac: 00:00:00:11:11:11
          [3] servId: 1
          [4] sapId: 1/1/1:111
   [53] Message type: Ack
   [54] DHCP server addr: 10.11.11.1
   [51] Lease time: 43200
   [1] Subnet mask: 255.255.255.0
   [3] Router: 10.1.1.251
   [6] Domain name server: length = 8
            2.2.2.2
            2.2.2.1
   [15] Domain name: domain.org
   [44] NETBIOS name server: 10.1.1.252
   [46] NETBIOS type: 1
   [251] Unknown option: len = 3, value = 01 02 03
   [254] Unknown option: len = 38, value = 07 06 73 75 62 2d 31 31 08 0d 73
   6c 61 2d 70 72 6f 66 69 6c 65 2d 31 09 0d 73 75 62 2d 70 72 6f 66 69 6c 65
   2d 31
   [60] Class id: (hex) aa bb
   [255] End
"
```

The user-defined option value for carrying the ESM strings and taken from the LUDB (identification-strings) should match the value defined in the strings-from-option parameter (value 254 in the following example) referenced in the subscriber identification policy on the relaying node.

```
configure
    subscriber-mgmt
        sub-ident-policy "sub-id-pol-1" create
            sub-profile-map
                use-direct-map-as-default
            exit
            sla-profile-map
                use-direct-map-as-default
            exit
            strings-from-option 254
        exit
```

The DHCP options in the debug output match the definition of the options in the LUBD.

Figure 190 **Decoding the ESM User Option**

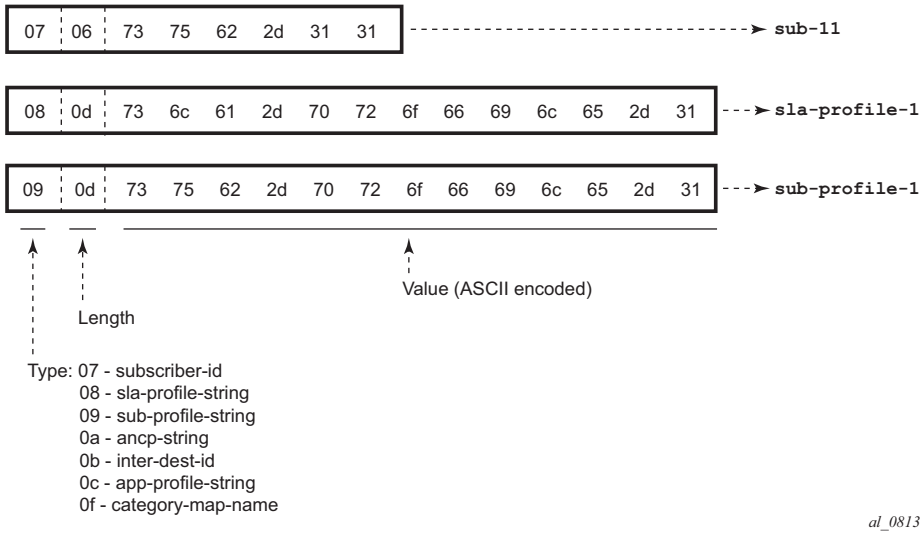

PPP Users Verification

The following command shows the DHCPv4 server lease state record for LUDB host entry-55. The client type is set to ppp, and the address type is set to dynamic as ludb-1 returns address-pool pool4-2 for this host.

```
*A:BNG-1# show router dhcp local-dhcp-server "dhcp4" leases 10.1.2.2 detail
===============================================================================
Lease for DHCP server dhcp4 router Base
===============================================================================
IP-address                    : 10.1.2.2
Lease-state                   : stable
Lease started                 : 2015/05/29 14:30:16
```

```
Last renew                    : N/A
Remaining LifeTime            : 0h51m9s
Remaining Potential Exp. Time: 0h0m0s
MAC address                   : 00:00:00:55:55:55
Xid                           : 0xeb633130
Failover Control              : local
Client Type                   : ppp
User-db Host Name             : entry-55
User-db Address Type          : dynamic
Persistence Key               : N/A
Opt82 Hex Dump                : (length=83)
                              : 52 51 01 1d 42 4e 47 2d 31 7c 31 7c 67 72 70 2d
                              : 69 6e 74 2d 31 2d 31 7c 31 2f 31 2f 31 3a 31 31
                              : 31 02 06 00 00 00 55 55 55 09 28 00 00 19 7f 23
                              : 02 06 00 00 00 55 55 55 06 01 01 01 05 42 4e 47
                              : 2d 31 03 04 00 00 00 01 04 09 31 2f 31 2f 31 3a
                              : 31 31 31
Opt82 Circuit Id              : BNG-1|1|grp-int-1-1|1/1/1:111
Opt82 Remote Id               : (hex) 00 00 00 55 55 55
Opt82 VS System               : BNG-1
Opt82 VS Clnt MAC             : 00:00:00:55:55:55
Opt82 VS Service              : (hex) 00 00 00 01
Opt82 VS SAP                  : 1/1/1:111
Opt82 VS String               :
Opt82 VS PPPoE Session ID     :
Opt60 Hex Dump                : (length=10)
                              : 41 4c 55 37 58 58 58 53 42 4d
Lease Remaining Hold Time     : 0h0m0s
===============================================================================
*A:BNG-1#
```

The debug output on the DHCPv4 server and the LUDB ludb-1 shows that the LUDB is accessed for every incoming message. Again the identification-strings are returned to the relay agent in the Offer and Acknowledge messages using option [254].

```
79 2015/06/05 13:55:09.40 UTC MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4
Rx DHCP Discover

   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 10.1.1.254
   chaddr: 00:00:00:55:55:55    xid: 0xbe277b30

   DHCP options:
   [82] Relay agent information: len = 87
      [1] Circuit-id: BNG-1|1|grp-int-1-1|1/1/1:111
      [2] Remote-id: (hex) 00 00 00 55 55 55
      [9] Vendor-Specific info: len = 46
          Enterprise [6527] : len = 41
          [2] clntMac: 00:00:00:55:55:55
          [6] clntType: 1
          [1] systemId: BNG-1
          [3] servId: 1
          [4] sapId: 1/1/1:111
          [17] pppoeSessionId: 1
   [51] Lease time: 3600
```

```
   [53] Message type: Discover
   [60] Class id: ALU7XXXSBM
   [255] End
"

80 2015/06/05 13:55:09.40 UTC MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  mac: 00:00:00:55:55:55

  Host entry-55 found in user data base ludb-1"

81 2015/06/05 13:55:09.40 UTC MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4
lease added for 10.1.2.20 state=offer
"

82 2015/06/05 13:55:09.41 UTC MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4
Tx DHCP Offer to local client 10.1.1.254 vrId=1

   ciaddr: 0.0.0.0           yiaddr: 10.1.2.20
   siaddr: 10.11.11.1        giaddr: 10.1.1.254
   chaddr: 00:00:00:55:55:55    xid: 0xbe277b30

   DHCP options:
   [82] Relay agent information: len = 87
      [1] Circuit-id: BNG-1|1|grp-int-1-1|1/1/1:111
      [2] Remote-id: (hex) 00 00 00 55 55 55
      [9] Vendor-Specific info: len = 46
          Enterprise [6527] : len = 41
          [2] clntMac: 00:00:00:55:55:55
          [6] clntType: 1
          [1] systemId: BNG-1
          [3] servId: 1
          [4] sapId: 1/1/1:111
          [17] pppoeSessionId: 1
   [53] Message type: Offer
   [54] DHCP server addr: 10.11.11.1
   [51] Lease time: 3600
   [1] Subnet mask: 255.255.255.0
   [6] Domain name server: 2.2.2.2
   [254] Unknown option: len = 32, value = 07 06 73 75 62 2d 35 35 08 0a 73
   6c 61 2d 70 72 6f 66 2d 35 09 0a 73 75 62 2d 70 72 6f 66 2d 33
   [60] Class id: ALU7XXXSBM
   [255] End
"

83 2015/06/05 13:55:09.41 UTC MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4
Rx DHCP Request

   ciaddr: 10.1.2.20         yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 10.1.1.254
   chaddr: 00:00:00:55:55:55    xid: 0xbe277b30

   DHCP options:
   [82] Relay agent information: len = 87
      [1] Circuit-id: BNG-1|1|grp-int-1-1|1/1/1:111
      [2] Remote-id: (hex) 00 00 00 55 55 55
```

```
      [9] Vendor-Specific info: len = 46
          Enterprise [6527] : len = 41
          [2] clntMac: 00:00:00:55:55:55
          [6] clntType: 1
          [1] systemId: BNG-1
          [3] servId: 1
          [4] sapId: 1/1/1:111
          [17] pppoeSessionId: 1
   [50] Requested IP addr: 10.1.2.20
   [51] Lease time: 3600
   [53] Message type: Request
   [54] DHCP server addr: 10.11.11.1
   [60] Class id: ALU7XXXSBM
   [255] End
"

84 2015/06/05 13:55:09.41 UTC MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  mac: 00:00:00:55:55:55

  Host entry-55 found in user data base ludb-1"

85 2015/06/05 13:55:09.41 UTC MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4
lease update for 10.1.2.20 state=stable
"

86 2015/06/05 13:55:09.41 UTC MINOR: DEBUG #2001 Base DHCP server
"DHCP server:  dhcp4
Tx DHCP Ack to local client 10.1.1.254 vrId=1

   ciaddr: 10.1.2.20         yiaddr: 10.1.2.20
   siaddr: 10.11.11.1        giaddr: 10.1.1.254
   chaddr: 00:00:00:55:55:55    xid: 0xbe277b30

   DHCP options:
   [82] Relay agent information: len = 87
      [1] Circuit-id: BNG-1|1|grp-int-1-1|1/1/1:111
      [2] Remote-id: (hex) 00 00 00 55 55 55
      [9] Vendor-Specific info: len = 46
          Enterprise [6527] : len = 41
          [2] clntMac: 00:00:00:55:55:55
          [6] clntType: 1
          [1] systemId: BNG-1
          [3] servId: 1
          [4] sapId: 1/1/1:111
          [17] pppoeSessionId: 1
   [53] Message type: Ack
   [54] DHCP server addr: 10.11.11.1
   [51] Lease time: 3600
   [1] Subnet mask: 255.255.255.0
   [6] Domain name server: 2.2.2.2
   [254] Unknown option: len = 32, value = 07 06 73 75 62 2d 35 35 08 0a 73
   6c 61 2d 70 72 6f 66 2d 35 09 0a 73 75 62 2d 70 72 6f 66 2d 33
   [60] Class id: ALU7XXXSBM
   [255] End
"
```

Operational Considerations

A DHCPv4 server with an LUDB cannot be used for supporting Local Address Assignment (LAA) scenarios. LAA can be used when the DHCPv4 relay agent and the DHCPv4 server are in the same node, and where IP address assignment for PPP users happens through the API directly into the local DHCPv4 server. See the ESM SLAAC Prefix Assignment via Local Address Server chapter via Local Address Server of the Advanced Configuration Guide for details on this topic.

Conclusion

This chapter explained and demonstrated the use of an LUDB in combination with a DHCPv4 server. For the DHCPv4 server to find matching entries in the LUDB, the relay agent has to be configured with the correct options and sub-options so that entries can be matched in the LUDB. It was noted that the DHCPv4 server does not require the relaying function to be located in the same node. The input and the output parameters of a DHCPv4 server-attached LUDB were listed. Carrying over the identification/ESM strings in a user-defined DHCP option was configured and demonstrated, and the decoding of this option was explained.

Local User Database for Enhanced Subscriber Management

Applicability

This chapter is applicable to the 7x50 SR series and was tested on SR-OS 13.0.R6.

Having knowledge of ESM Basics, the Routed CO model, and Local User Database Basics are prerequisites for understanding this note.

Summary

A Local User Database (LUDB) is a data source providing Enhanced Subscriber Management (ESM) data so that subscribers and subscriber hosts can be instantiated when end-users connect their devices. ESM data includes identification strings, IP address/prefix, profiles, and so on. See the ESM Basics chapter for more information.

LUDBs offer a self-contained method for providing the ESM data, so that no additional ESM data sources are needed.

Alternative ESM data sources are: RADIUS, Diameter NASREQ, Diameter Gx, DHCP-server, Python, and defaults.

Mixed scenarios, where part of the data is provided by an LUDB and the remaining part is provided through RADIUS, are the subject of the Flexible Authentication Model in ESM chapter.

LUDBs can be used for the following applications:

- assisting a DHCPv4 server in assigning fixed IP addresses to dedicated devices; see the Local User Database for DHCPv4 Server chapter.
- authenticating devices, so that ESM hosts and subscribers can be instantiated. This is supported for the Routed CO model only.
- authenticating devices as a fallback for RADIUS authentication, in case the RADIUS server is not available. This is supported for the Routed CO model only.

This chapter describes the use of LUDBs for authentication, including:

- parameters that can be returned by LUDBs
- contexts where LUDBs can be applied in the system

Overview

LUDB authentication is supported:

- for IPoE as well as for PPP
- for regular SAPs as well as for capture and managed SAPs
- for the proxy scenario as well as for the relay scenario

LUDB authentication can be started directly through one of the following protocol triggers; see Figure 191:

- DHCPv4 Discover
- DHCPv6 Solicit
- PPPoE PADI
- PPPoA Conf Req
- Router Solicit [RS]

Figure 191 **LUDB Authentication**

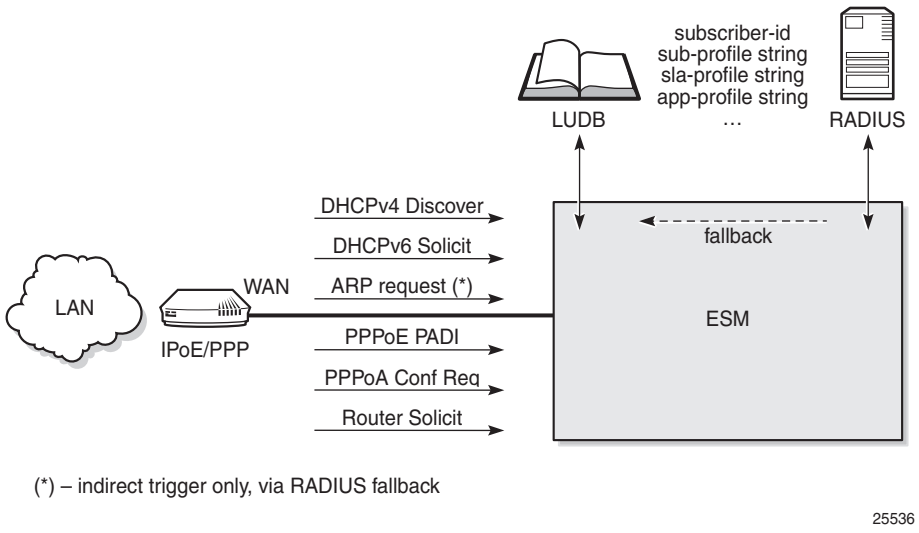

When triggered, ESM can directly access an LUDB because the LUDB is applied to the service directly (through one of its sub-contexts), or indirectly as a fallback action for RADIUS (through the authentication policy); see Figure 192. ARP requests can only trigger LUDB authentication indirectly.

Note: An authentication policy can be referenced from a group interface and a capture SAP, or from an LUDB.

Figure 192 **Direct and Indirect LUDB Authentication**

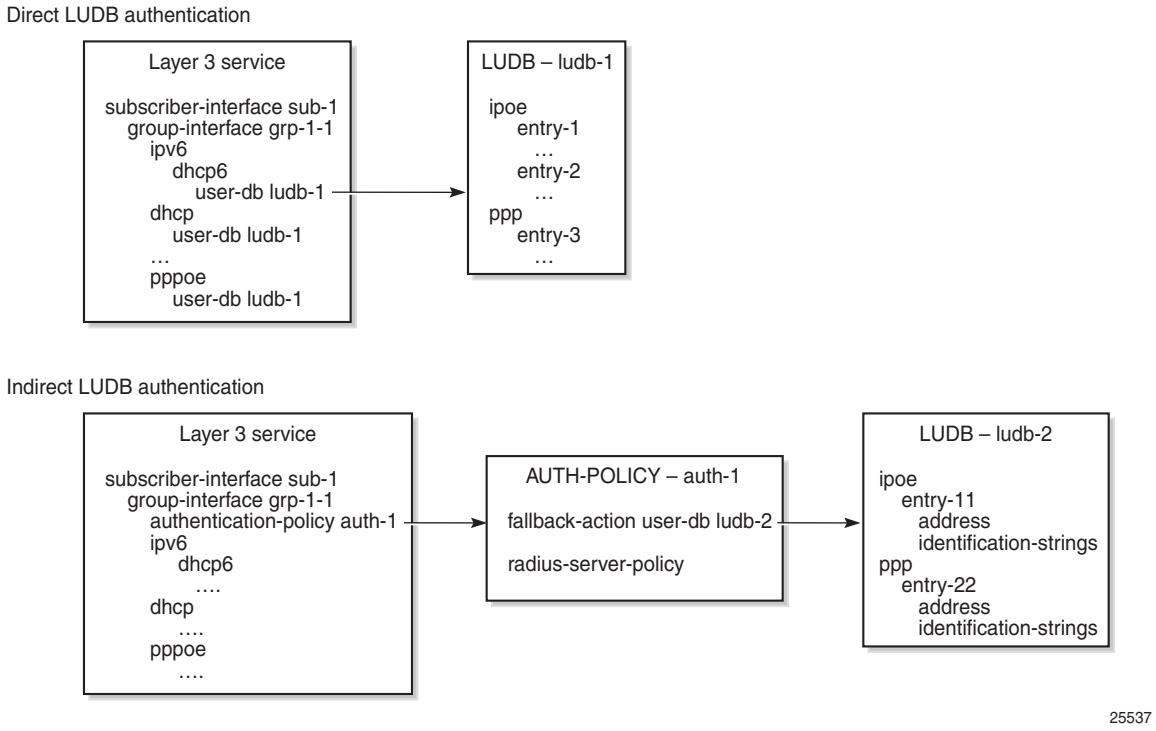

Three ESM scenarios in which an LUDB is accessed are as follows:

- ESM gets all the data needed for host creation directly from the LUDB.
- ESM gets some data from the LUDB and the remaining data from an AAA server (RADIUS, NASREQ, Gx). This requires the LUDB to provide an authentication or a Diameter application policy, and no authentication or Diameter policy at the group interface level.
- ESM tries to fetch the data from a RADIUS server, but because this server is not reachable, ESM falls back to an LUDB.

The examples in the Configuration section of this chapter describe the first and the last scenario. The second scenario is described in the Flexible Authentication Model in ESM chapter.

LUDB Input and Output Parameters

As described in the Local User Database Basics chapter, when processing an LUDB lookup request, the input parameters are filtered and optionally masked before searching through the entries in the database. Every entry, except for the default, contains one or more host-identification fields that are used for matching purposes. As a result of the lookup process, these output parameters are then used for host creation.

The following IPoE host-identification fields are supported when accessing an LUDB for ESM; see Figure 193:

- mac
- circuit-id + remote-id
- option60 (excluded for the IPoE session model)
- sap-id + encap-tag-range
- service-id (7x50 service-id)
- string
- system-id

The LUDB lookup process can take up to four IPoE match-criteria into account, as defined by the IPoE match-list.

The following PPP host-identification fields are supported when accessing an LUDB for ESM; see Figure 194:

- mac
- circuit-id + remote-id
- sap-id + encap-tag-range
- service-name (PPPoE tag: service name)
- username (excluded for the RADIUS fallback scenario)

The LUDB lookup process can take up to three PPP match-criteria into account, as defined by the PPP match-list.

The fields output from the lookup process include the identification strings, options, and others.

See Figure 193 and Figure 194 for the full list of input and output parameters for IPoE and PPPoE, respectively.

Figure 193 **LUDB parameters for IPoE**

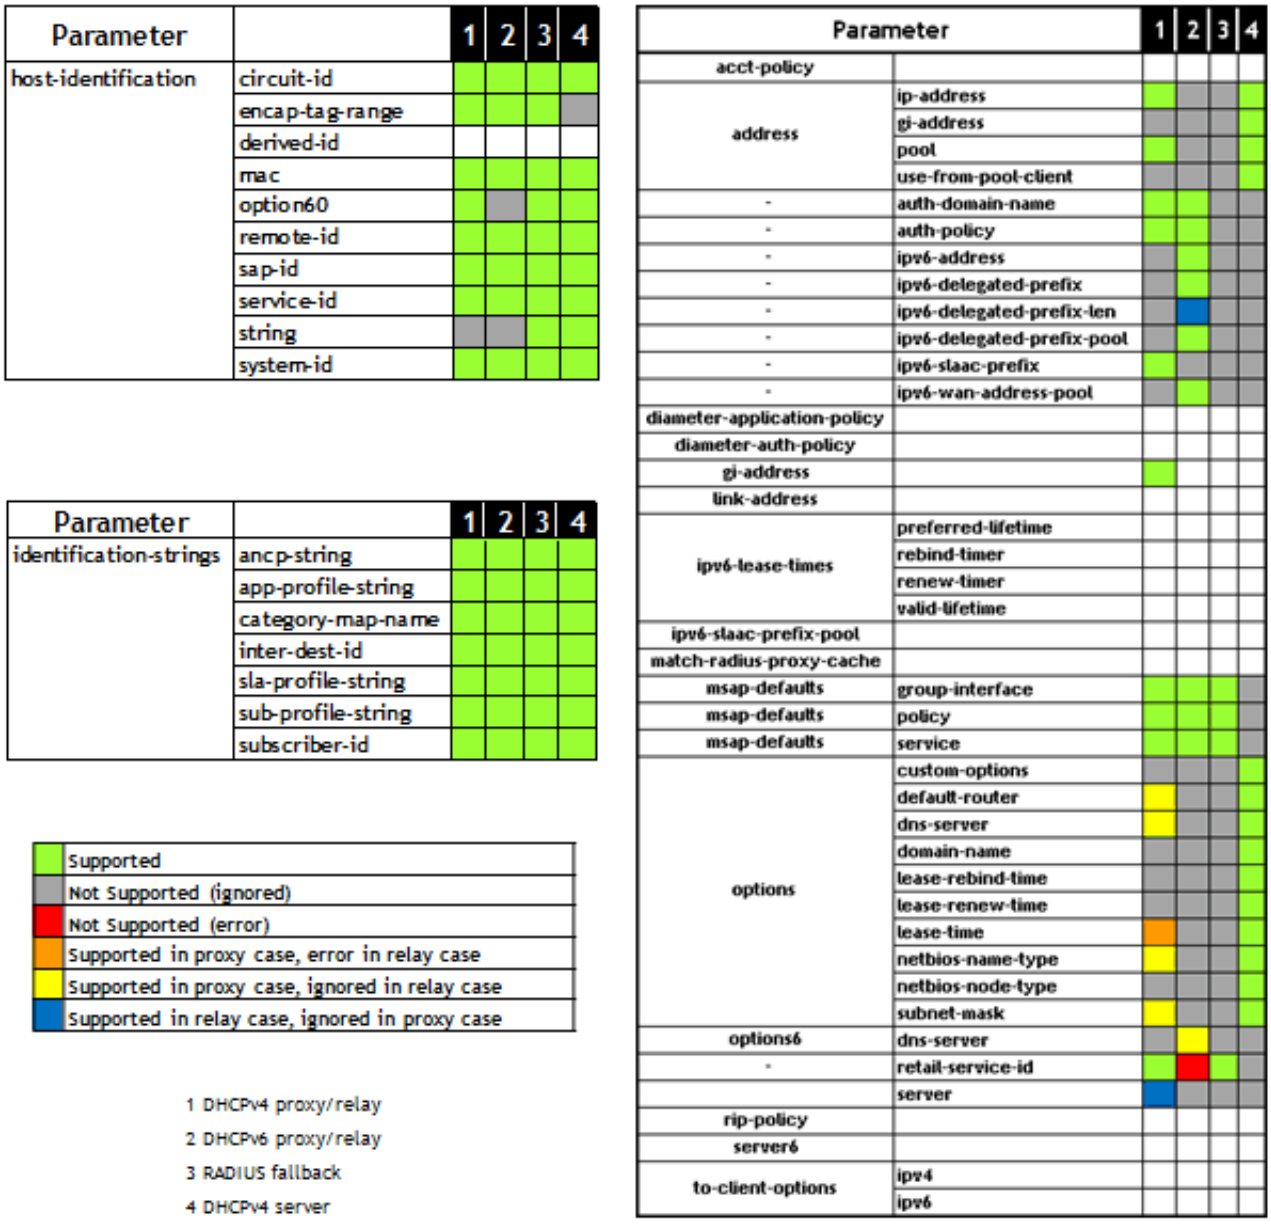

Parameter		1	2	3	4
host-identification	circuit-id				
	encap-tag-range				
	derived-id				
	mac				
	option60				
	remote-id				
	sap-id				
	service-id				
	string				
	system-id				

Parameter		1	2	3	4
identification-strings	ancp-string				
	app-profile-string				
	category-map-name				
	inter-dest-id				
	sla-profile-string				
	sub-profile-string				
	subscriber-id				

	Supported
	Not Supported (ignored)
	Not Supported (error)
	Supported in proxy case, error in relay case
	Supported in proxy case, ignored in relay case
	Supported in relay case, ignored in proxy case

1 DHCPv4 proxy/relay
2 DHCPv6 proxy/relay
3 RADIUS fallback
4 DHCPv4 server

Parameter		1	2	3	4
acct-policy					
address	ip-address				
	gi-address				
	pool				
	use-from-pool-client				
-	auth-domain-name				
-	auth-policy				
-	ipv6-address				
-	ipv6-delegated-prefix				
-	ipv6-delegated-prefix-len				
-	ipv6-delegated-prefix-pool				
-	ipv6-slaac-prefix				
-	ipv6-wan-address-pool				
diameter-application-policy					
diameter-auth-policy					
gi-address					
link-address					
ipv6-lease-times	preferred-lifetime				
	rebind-timer				
	renew-timer				
	valid-lifetime				
ipv6-slaac-prefix-pool					
match-radius-proxy-cache					
msap-defaults	group-interface				
msap-defaults	policy				
msap-defaults	service				
options	custom-options				
	default-router				
	dns-server				
	domain-name				
	lease-rebind-time				
	lease-renew-time				
	lease-time				
	netbios-name-type				
	netbios-node-type				
	subnet-mask				
options6	dns-server				
-	retail-service-id				
	server				
rip-policy					
server6					
to-client-options	ipv4				
	ipv6				

Figure 194 **LUDB parameters for PPPoE**

Parameter		1	2	3
host-identification	circuit-id			
	encap-tag-range			
	mac			
	remote-id			
	sap-id			
	service-name			
	username			

Parameter		1	2	3
identification-strings	ancp-string			
	app-profile-string			
	category-map-name			
	inter-dest-id			
	sla-profile-string			
	sub-profile-string			
	subscriber-id			

	Supported
	Not Supported (ignored)
	Not Supported (error)
	Supported in proxy case, error in relay case
	Supported in proxy case, ignored in relay case
	Supported in relay case, ignored in proxy case

1 DHCPv4/6 PROXY / RELAY
2 RADIUS FALLBACK
3 DHCPv4 SERVER

Parameter		1	2	3
access-loop-encapsulation	encap-offset			
	rate-down			
access-loop-information	circuit-id			
	remote-id			
acct-policy				
address	ip-address			
	gi-address			
	pool			
	use-from-pool-client			
-	auth-policy			
	force-ipv6cp			
diameter-application-policy				
diameter-auth-policy				
ignore-df-bit				
ipv6-lease-times	preferred-lifetime			
	rebind-timer			
	renew-timer			
	valid-lifetime			
ipv6-slaac-prefix-pool				
interface	service-id			
-	ipv6-address			
-	ipv6-delegated-prefix			
-	ipv6-delegated-prefix-len			
-	ipv6-delegated-prefix-pool			
-	ipv6-slaac-prefix			
-	ipv6-wan-address-pool			
l2tp	group service-id			
msap-defaults	group-interface			
msap-defaults	policy			
msap-defaults	service			
options	custom-options			
	dns-server			
	netbios-name-server			
options6	dns-server			
	pado-delay			
	password			
-	pre-auth-policy			
	retail-service-id			
rip-policy				
to-client-options	ipv4			
	ipv6			

Applying an LUDB for ESM

LUDB authentication for regular SAPs requires an LUDB to be applied at the group interface level in the Layer 3 service (VPRN or IES); see Figure 195. All the SAPs on that group interface share the same authentication configuration. See the Local User Database for DHCPv4 Server chapter for the scenario where a user database is attached to a DHCPv4 server.

Figure 195 **LUDB Authentication for Regular SAPs**

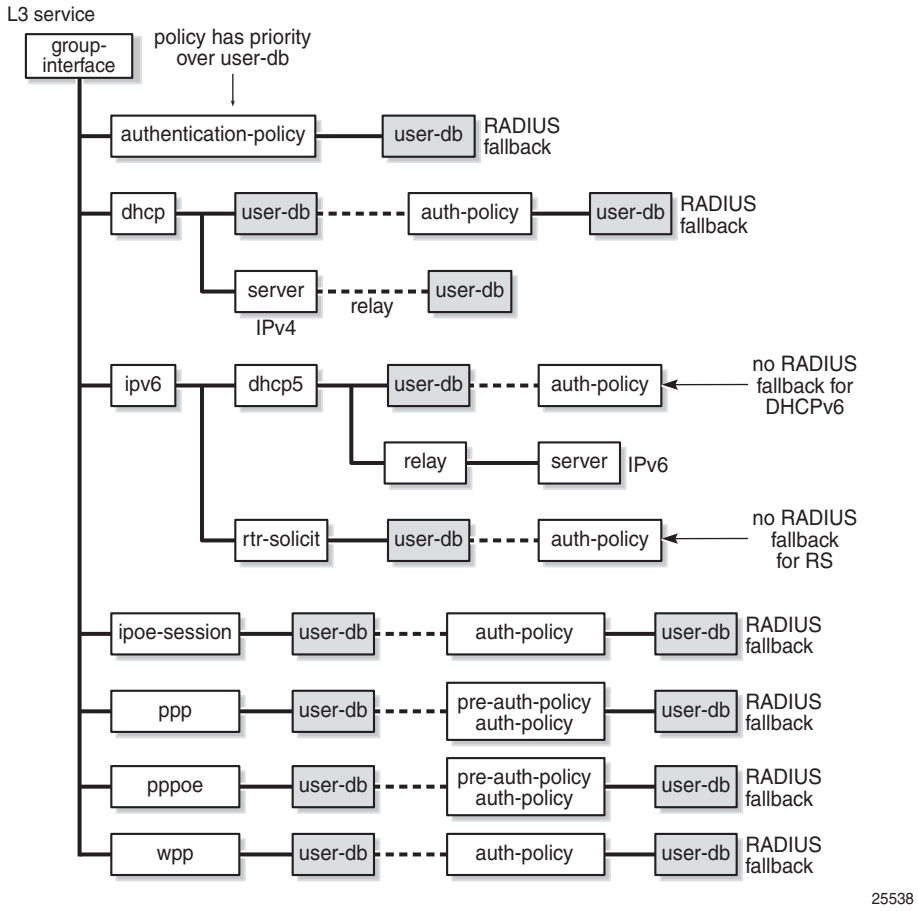

LUDB authentication for capture and managed SAPs requires an LUDB to be assigned at capture SAP level in the Layer 2 service (capture-VPLS), and at the group interface level in the Layer 3 service (VPRN or IES). Because the trigger messages to create the managed SAPs are received on the capture SAP and subsequent messages on the managed SAP, the authentication configurations for the Layer 2 and the Layer 3 service must align, including the LUDBs; see Figure 196.

Figure 196 **LUDB Authentication for Capture and Managed SAPs**

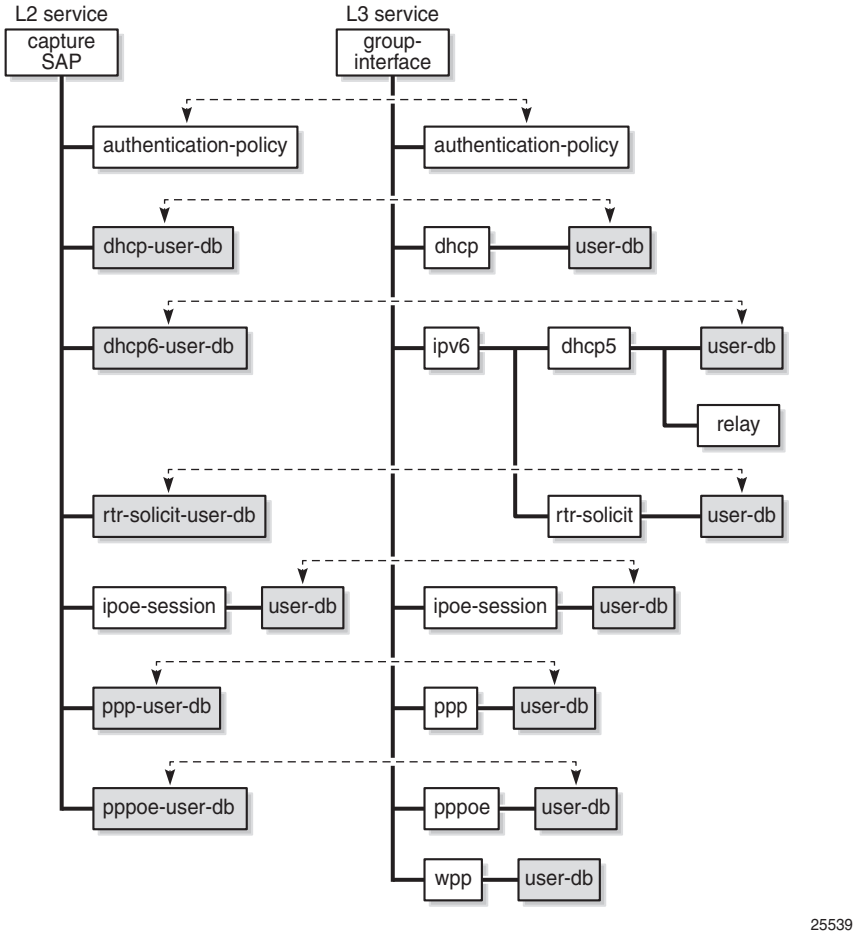

The following CLI commands are available for applying LUDBs:

```
configure service vprn|ies subscriber-interface group-
 interface
     dhcp user-db <local-user-db-name>
     ipoe-session user-db <local-user-db-name>
     ipv6 dhcp6 user-db <local-user-db-name>
     ipv6 router-solicit user-db <local-user-db-name>
     ppp user-db <local-user-db-name>
     pppoe user-db <local-user-db-name>
     wpp user-db <local-user-db-name>

configure service vpls sap
     dhcp-user-db <local-user-db-name>
     dhcp6-user-db <local-user-db-name>
```

```
ipoe-session user-db <local-user-db-name>
ppp-user-db <local-user-db-name>
pppoe-user-db <local-user-db-name>
rtr-solicit-user-db <local-user-db-name>
```

An LUDB can be assigned in different contexts, and can be reused. Assuming an LUDB contains both IPoE as well as PPP entries, this LUDB is likely to be assigned in a dhcp context as well as in a ppp or a pppoe context.

Configuration Guidelines

The following rules have to be observed when configuring authentication for regular, capture, and managed SAPs:

- If an authentication policy is applied at the capture SAP or group interface level, that authentication policy has priority, no matter whether or in which other sub-contexts an LUDB is assigned. Only when the AAA/RADIUS server referenced from the authentication policy is not available, can the SR OS rely on a fallback LUDB if configured. In that case, only a limited set of parameters are returned; see Figure 193 and Figure 194.

  This means that for an LUDB to provide ESM data, no authentication policy may be applied at the capture SAP or group interface level, provided that the LUDB is in the no shutdown state.
- An LUDB can return an authentication policy so that the ESM data can be partially provided by the LUDB, and partially by an AAA/RADIUS server. For this mixed scenario, RADIUS fallback is only possible for PPP, PPPoE, DHCPv4, IPoE sessions, and WPP, but not for DHCPv6 and IPv6 router solicitation. For more information, see the Flexible Authentication Model in ESM chapter. When the AAA/RADIUS server is defined but not available, the SR OS can rely on a fallback LUDB if configured.
- LUDB authentication for RADIUS fallback requires an LUDB to be applied to an authentication policy as a fallback action:

```
configure subscriber-mgmt
    authentication-policy fallback-action user-db <local-user-db-name>
```

- The DHCPv4 server referenced from a group interface in the dhcp context (for supporting the relay scenario) can have an LUDB assigned; see the LUDB for DHCPv4 chapter. See Figure 193 and Figure 194 for the parameters that this LUDB can return to the DHCPv4 server.

Note: An LUDB cannot be assigned to a DHCPv6 server.

- If an LUDB is applied in the ipoe-session context of a group interface or capture SAP, the LUDBs assigned in the dhcp, dhcp6, and router-solicit contexts of the same group interface or capture SAP are ignored.

  This avoids accessing the LUDB on every DHCPv4 DORA or DHCPv6 SARR message, which is the case when no IPoE sessions are used.

  For IPoE sessions, the LUDB host identification cannot be based on option60. Entries in the LUDB with host-identification option60 strings are ignored. All the other LUDB entry match criteria are allowed.
- If an LUDB is applied in the ppp or pppoe context of a group interface or capture SAP, PAP or CHAP authentication is based on the password configured in the entry. If no password is required, the password parameter in the LUDB entry must be explicitly set to ignore.

  A password verification failure leads to a setup failure.

Configuration

Figure 197 shows the baseline configuration used in this chapter. Dual and single stack end-user devices supporting IPoE and PPPoE connect to the SAPs of IES-1. Different LUDBs are added to this baseline configuration later in this chapter, depending on the scenario.

Figure 197 **Baseline setup**

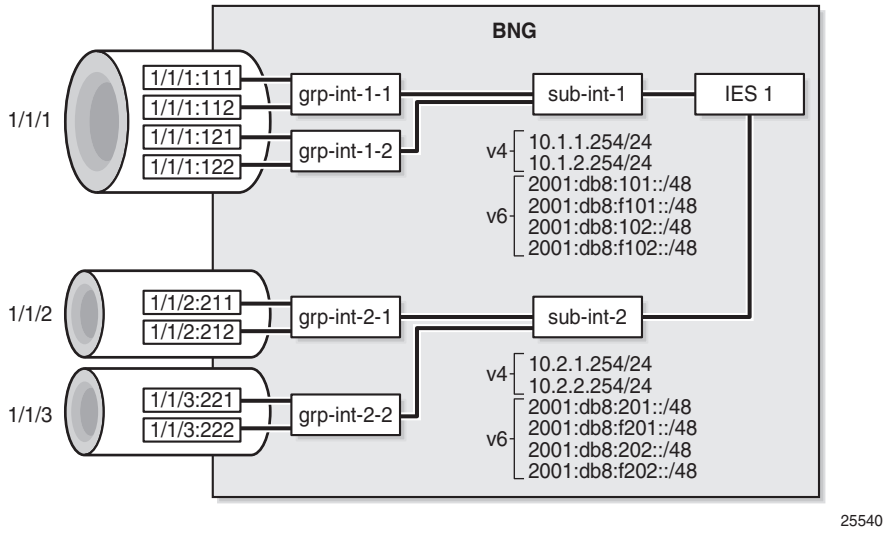

The following partial configuration applies to IES-1. This service is provisioned with ESM enabled on all of its SAPs, and supports proxy and relay scenarios on all group interfaces for both IPv4 and IPv6. Only the part relevant to subscriber interface "sub-int-1" and group interface "grp-int-1-1" is shown. The configurations for the other subscriber and group interfaces are similar. Check the ESM Basics and Routed CO chapters for more information.

```
configure
    service
        ies 1 customer 1 create
            subscriber-interface "sub-int-1" create
                address 10.1.1.254/24
                --- snipped ---
                ipv6
                    delegated-prefix-len 56
                    link-local-address fe80::ea:4b:f1
                    subscriber-prefixes
                        prefix 2001:db8:101::/48 wan-host
                        prefix 2001:db8:f101::/48 pd
                        --- snipped ---
                    exit
                exit
                group-interface "grp-int-1-1" create
                    ipv6
                        router-advertisements
                            no shutdown
                        exit
                        dhcp6
                            proxy-server
                                client-applications dhcp ppp
                                no shutdown
                            exit
```

```
                            relay
                                link-address 2001:db8:101::1
                                server 2001:db8::11
                                client-applications dhcp ppp
                                no shutdown
                            exit
                        exit
                        router-solicit
                            no shutdown
                        exit
                    exit
                    arp-populate
                    dhcp
                        proxy-server
                            emulated-server 10.1.1.254
                            no shutdown
                        exit
                        option
                            action keep
                        exit
                        server 10.11.11.1
                        trusted
                        lease-populate 100
                        client-applications dhcp ppp
                        gi-address 10.1.1.254
                        no shutdown
                    exit
                    sap 1/1/1:111 create
                        sub-sla-mgmt
                            def-sub-profile "sub-prof-1"
                            def-sla-profile "sla-prof-1"
                            sub-ident-policy "sub-id-pol-1"
                            multi-sub-sap
                            no shutdown
                        exit
                    exit
                    --- snipped ---
                exit
            exit
            --- snipped ---
```

For brevity, the configurations of the local DHCPv4 and DHCPv6 servers are not shown.

An excerpt from the LUDB "ludb-rsap" follows. Host "entry-11" defines the settings for a dual stack IPoE host, and host "entry-55" the settings for a dual stack PPPoE host. For both hosts, the LUDB provides all the data needed to ensure host instantiation.

```
configure
    subscriber-mgmt
        local-user-db "ludb-rsap" create
            description "LUDB for regular SAPs"
            ipoe
                match-list mac
                host "entry-11" create
                    host-identification
```

```
                    mac 00:00:00:11:11:11
                exit
                address 10.1.1.211
                identification-strings 254 create
                    subscriber-id "sub-11"
                    sla-profile-string "sla-prof-1"
                    sub-profile-string "sub-prof-1"
                exit
                options
                    subnet-mask 255.255.255.0
                    default-router 10.1.1.254
                    dns-server 2.2.2.2 2.2.2.1
                    domain-name "domain.org"
                    custom-option 251 hex 0x010203
                exit
                options6
                    dns-server 2001:db8:ddd:1::1 2001:db8:ddd:2::1
                exit
                ipv6-address 2001:db8:102:11::11
                ipv6-delegated-prefix 2001:db8:f102:1100::/56
                ipv6-delegated-prefix-len 56
                no shutdown
            exit
            --- snipped ---
        exit
        ppp
            match-list username
            host "entry-55" create
                host-identification
                    username "sub55@domain1"
                exit
                address 10.1.1.225/24
                password chap sub55
                identification-strings 254 create
                    subscriber-id "sub-55"
                    sla-profile-string "sla-prof-5"
                    sub-profile-string "sub-prof-3"
                exit
                options
                    dns-server 2.2.2.2
                exit
                options6
                    dns-server 2001:db8:ddd:1::1 2001:db8:ddd:2::1
                exit
                ipv6-address 2001:db8:101:55::55
                ipv6-delegated-prefix 2001:db8:f101:5500::/56
                ipv6-delegated-prefix-len 56
                no shutdown
            exit
            --- snipped ---
        exit
        no shutdown
    exit
```

IPoE Authentication - Session Model

In this example, the LUDB "ludb-rsap" is applied to the group interface in the ipoe-session context. This is the Nokia recommended way for supporting IPoE subscribers through an LUDB.

```
configure
    service
        ies 1 customer 1 create
            subscriber-interface "sub-int-1"
                group-interface "grp-int-1-1"
                    ipoe-session
                        ipoe-session-policy "ipoe-sess-1"
                        session-limit 100
                        user-db "ludb-rsap"
                        no shutdown
                    exit
```

Use the following debug configuration for troubleshooting connection issues.

```
debug
    router "Base"
        ip
            dhcp
                detail-level low
                mode egr-ingr-and-dropped
            exit
            dhcp6
                mode egr-ingr-and-dropped
                detail-level low
            exit
        exit
    exit
    subscriber-mgmt
        local-user-db "ludb-rsap"
            detail all
        exit
    exit
exit
```

The following trace appears when the user with MAC address 00:00:00:11:11:11 first connects using DHCPv4 and subsequently connects using DHCPv6 without removing the DHCPv4 connection.

Note: The LUDB is accessed just once, immediately after the DHCPv4 Discover message.

```
119 2015/10/30 15:44:08.12 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 4 (grp-int-1-1),
   received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:111) Port 67
```

```
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:00:00:11:11:11    xid: 0x6
"

120 2015/10/30 15:44:08.12 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  mac: 00:00:00:11:11:11
  Host entry-11 found in user data base ludb-rsap"

121 2015/10/30 15:44:08.12 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 4 (grp-int-1-1),
   transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:111) Port 68
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.1.1.211
   siaddr: 10.1.1.254        giaddr: 10.1.1.254
   chaddr: 00:00:00:11:11:11    xid: 0x6
"

122 2015/10/30 15:44:08.16 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 4 (grp-int-1-1),
   received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:111) Port 67
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:00:00:11:11:11    xid: 0x6
"

123 2015/10/30 15:44:08.16 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 4 (grp-int-1-1),
   transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:111) Port 68
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.1.1.211
   siaddr: 10.1.1.254        giaddr: 10.1.1.254
   chaddr: 00:00:00:11:11:11    xid: 0x6
"
124 2015/10/30 15:44:23.22 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Incoming DHCP6 Msg : SOLICIT (1)
   on itf grp-int-1-1
"

125 2015/10/30 15:44:23.22 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Outgoing DHCP6 Msg : ADVERTISE (2)
   to itf grp-int-1-1
"

126 2015/10/30 15:44:23.23 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Incoming DHCP6 Msg : REQUEST (3)
   on itf grp-int-1-1
"

127 2015/10/30 15:44:23.23 CET MINOR: DEBUG #2001 Base TIP
```

```
"TIP: DHCP6_PKT
   Outgoing DHCP6 Msg : REPLY (7)
   to itf grp-int-1-1
"
The active subscriber hosts for which the address origin is ludb are shown with the
following command.
*A:BNG-1# show service id 1 subscriber-hosts address-origin ludb
===============================================================
Subscriber Host table
===============================================================
Sap                    Subscriber
  IP Address
    MAC Address           PPPoE-SID Origin       Fwding State
---------------------------------------------------------------
1/1/1:111              sub-11
  10.1.1.211
    00:00:00:11:11:11    N/A       DHCP         Fwding
1/1/1:111              sub-11
  2001:db8:102:11::11/128
    00:00:00:11:11:11    N/A       IPoE-DHCP6   Fwding
1/1/1:111              sub-11
  2001:db8:f102:1100::/56
    00:00:00:11:11:11    N/A       IPoE-DHCP6   Fwding
---------------------------------------------------------------
Number of subscriber hosts : 3
===============================================================
*A:BNG-1#
```

The following command shows the session details for MAC address 00:00:00:11:11:11.

Note: This information aligns with the LUDB configuration of "ludb-rsap", and the origin codes are set to UserDb.

```
*A:BNG-1# show service id 1 ipoe session mac 00:00:00:11:11:11 detail
===============================================================================
IPoE sessions for service 1
===============================================================================
SAP                     : 1/1/1:111
Mac Address             : 00:00:00:11:11:11
Circuit-Id              : 11
Remote-Id               : AA
Session Key             : sap-mac
MC-Standby              : No
Subscriber-interface    : sub-int-1
Group-interface         : grp-int-1-1
Termination Type        : local
Up Time                 : 0d 00:00:44
Session Time Left       : N/A
Last Auth Time          : 10/30/2015 15:44:09
Min Auth Intvl (left)   : infinite (N/A)
Persistence Key         : N/A
Subscriber              : "sub-11"
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         : "EA4BFF0000338256338238"
Sap-Session-Index       : 1
IP Address              : 10.1.1.211/24
IP Origin               : UserDb
Primary DNS             : 2.2.2.2
Press any key to continue (Q to quit)

Secondary DNS           : 2.2.2.1
Primary NBNS            : N/A
Secondary NBNS          : N/A
Address-Pool            : N/A
IPv6 Prefix             : N/A
IPv6 Prefix Origin      : None
IPv6 Prefix Pool        : ""
IPv6 Del.Pfx.           : 2001:db8:f102:1100::/56
IPv6 Del.Pfx. Origin    : UserDb
IPv6 Del.Pfx. Pool      : ""
IPv6 Address            : 2001:db8:102:11::11
IPv6 Address Origin     : UserDb
IPv6 Address Pool       : ""
Primary IPv6 DNS        : 2001:db8:ddd:1::1
Secondary IPv6 DNS      : 2001:db8:ddd:2::1
Radius Session-TO       : N/A
Radius Class            :
Radius User-Name        :
-------------------------------------------------------------------------------
Number of sessions : 1
===============================================================================
*A:BNG-1#
```

Other commands displaying origin codes are:

```
*A:BNG-1# show service id 1 dhcp lease-state mac 00:00:00:11:11:11
===============================================================================
DHCP lease state table, service 1
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LeaseTime  Origin   Stdby
-------------------------------------------------------------------------------
10.1.1.211      00:00:00:11:11:11 1/1/1:111           06d23h59m  UserDb
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
*A:BNG-1#
*A:BNG-1# show service id 1 dhcp6 lease-state mac 00:00:00:11:11:11
===============================================================================
DHCP lease state table, service 1
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LeaseTime  Origin   Stdby
-------------------------------------------------------------------------------
2001:db8:102:11::11/128
                00:00:00:11:11:11 1/1/1:111           23h59m12s  UserDb
2001:db8:f102:1100::/56
```

```
               00:00:00:11:11:11 1/1/1:111           23h59m12s  UserDb
-------------------------------------------------------------------------------
Number of lease states : 2
===============================================================================
*A:BNG-1#
```

IPoE Authentication - Host Model

In this example, the LUDB "ludb-rsap" is applied to the group interface in the dhcp6, router-solicit, and dhcp contexts, but not in the ipoe-session context.

```
configure
    service
        ies 1
            subscriber-interface "sub-int-1"
                group-interface "grp-int-1-1"
                    ipv6
                        dhcp6
                            user-db "ludb-rsap"
                        exit
                        router-solicit
                            user-db "ludb-rsap"
                            no shutdown
                        exit
                    exit
                    dhcp
                        user-db "ludb-rsap"
                        no shutdown
                    exit
                exit
```

With the same debug configuration as for the IPoE session model, the LUDB is accessed multiple times when devices connect, as shown in the following trace.

```
131 2015/10/30 15:46:18.78 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 4 (grp-int-1-1),
   received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:111) Port 67
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:00:00:11:11:11    xid: 0x6
"

132 2015/10/30 15:46:18.78 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  mac: 00:00:00:11:11:11
  Host entry-11 found in user data base ludb-rsap"

133 2015/10/30 15:46:18.78 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 4 (grp-int-1-1),
   transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:111) Port 68
```

```
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.1.1.211
   siaddr: 10.1.1.254        giaddr: 10.1.1.254
   chaddr: 00:00:00:11:11:11    xid: 0x6
"

134 2015/10/30 15:46:18.80 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 4 (grp-int-1-1),
   received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:111) Port 67
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:00:00:11:11:11    xid: 0x6
"

135 2015/10/30 15:46:18.80 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  mac: 00:00:00:11:11:11
  Host entry-11 found in user data base ludb-rsap"

136 2015/10/30 15:46:18.80 CET MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 4 (grp-int-1-1),
   transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:111) Port 68
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.1.1.211
   siaddr: 10.1.1.254        giaddr: 10.1.1.254
   chaddr: 00:00:00:11:11:11    xid: 0x6
"
137 2015/10/30 15:46:44.46 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Incoming DHCP6 Msg : SOLICIT (1)
   on itf grp-int-1-1
"

138 2015/10/30 15:46:44.46 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  mac: 00:00:00:11:11:11
  Host entry-11 found in user data base ludb-rsap"

139 2015/10/30 15:46:44.46 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Outgoing DHCP6 Msg : ADVERTISE (2)
   to itf grp-int-1-1
"

140 2015/10/30 15:46:44.46 CET MINOR: DEBUG #2001 Base TIP
"TIP: DHCP6_PKT
   Incoming DHCP6 Msg : REQUEST (3)
   on itf grp-int-1-1
"

141 2015/10/30 15:46:44.46 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  mac: 00:00:00:11:11:11
  Host entry-11 found in user data base ludb-rsap"

142 2015/10/30 15:46:44.47 CET MINOR: DEBUG #2001 Base TIP
```

```
"TIP: DHCP6_PKT
   Outgoing DHCP6 Msg : REPLY (7)
   to itf grp-int-1-1
"
```

The LUDB is accessed for every incoming message. In a proxy case, the LUDB is accessed two times per host because the downstream messages (Offer and Reply for IPv4, Solicit and Reply for IPv6) are generated by ESM. In a relay case, where an IP address or an IP prefix is allocated by the DHCP server, the LUDB is accessed four times per host.

The command to list the active subscriber hosts is the same as for the IPoE session model, and is not repeated here. The same applies to the other commands providing origin codes.

PPPoE Authentication

In this example, the LUDB "ludb-rsap" is applied to the group interface in the pppoe context.

```
configure
    service
        ies 1
            subscriber-interface "sub-int-1"
                group-interface "grp-int-1-1"
                    pppoe
                        user-db "ludb-rsap"
                        no shutdown
                    exit
                exit
```

The following debug configuration applies for this example.

```
debug
    service
        id 1
            ppp
                packet
                    mode egr-ingr-and-dropped
                    detail-level high
                    discovery
                    ppp
                    dhcp-client
                exit
            exit
        exit
    exit
    subscriber-mgmt
        local-user-db "ludb-rsap"
            detail all
        exit
```

```
    exit
exit
```

The trace shows that the LUDB "ludb-rsap" is accessed once when user "sub55@domain1" connects. In this example, the LUDB access is in the middle of the CHAP authentication.

```
--- snipped ---
156 2015/10/30 15:49:23.22 CET MINOR: DEBUG #2001 Base PPPoE
"PPPoE: RX Packet
   IES 1, SAP 1/1/1:111
   DMAC: ea:4b:01:01:00:01
   SMAC: 00:00:00:55:55:55
   Ether Type: 0x8864 (Session)
   PPPoE Header:
   Version: 1                Type      : 1
   Code   : 0x00             Session-Id: 0x0001 (1)
   Length : 36
   PPP:
   Protocol  : 0xc223 (CHAP)
   Code      : 2 (Response)
   Identifier: 1             Length    : 34
   Value-Size: 16
   Value     : 75 61 3e 55 f7 89 4f 0d a5 bd 43 95 90 aa 69 d7
   Name      : "sub55@domain1"
   Hex Packet Dump:
   11 00 00 01 00 24 c2 23 02 01 00 22 10 75 61 3e 55 f7 89 4f 0d a5 bd 43 95
   90 aa 69 d7 73 75 62 35 35 40 64 6f 6d 61 69 6e 31 00 00 00 00
"

157 2015/10/30 15:49:23.22 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  user-name:
    original:  sub55@domain1
    masked:    sub55@domain1
  Host entry-55 found in user data base ludb-rsap"

158 2015/10/30 15:49:23.22 CET MINOR: DEBUG #2001 Base PPPoE
"PPPoE: TX Packet
   IES 1, SAP 1/1/1:111
   DMAC: 00:00:00:55:55:55
   SMAC: ea:4b:01:01:00:01
   Ether Type: 0x8864 (Session)
   PPPoE Header:
   Version: 1                Type      : 1
   Code   : 0x00             Session-Id: 0x0001 (1)
   Length : 33
   PPP:
   Protocol  : 0xc223 (CHAP)
   Code      : 3 (Success)
   Identifier: 1             Length    : 31
   Message: "CHAP authentication success"
   Hex Packet Dump:
   11 00 00 01 00 21 c2 23 03 01 00 1f 43 48 41 50 20 61 75 74 68 65 6e 74 69
   63 61 74 69 6f 6e 20 73 75 63 63 65 73 73
"
--- snipped ---
```

With this dual stack PPP user connected, the subscriber hosts created are:

```
*A:BNG-1# show service id 1 subscriber-hosts
===============================================================
Subscriber Host table
===============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID Origin       Fwding State
---------------------------------------------------------------
1/1/1:111              sub-55
  10.1.1.225
    00:00:00:55:55:55    1         IPCP         Fwding
1/1/1:111              sub-55
  2001:db8:101:55::55/128
    00:00:00:55:55:55    1         PPP-DHCP6    Fwding
1/1/1:111              sub-55
  2001:db8:f101:5500::/56
    00:00:00:55:55:55    1         PPP-DHCP6    Fwding
---------------------------------------------------------------
Number of subscriber hosts : 3
===============================================================
*A:BNG-1#
```

Detailed session information for PPP user "sub55@domain55" shows the origin codes.

```
*A:BNG-1# show service id 1 ppp session user-name "sub55@domain1" detail

===============================================================================
PPP sessions for service 1
===============================================================================

User-Name            : sub55@domain1

Description          : svc:1 sap:1/1/1:111 mac:00:00:00:55:55:55 sid:1
Up Time              : 0d 00:00:34
Type                 : oE
Termination          : local
IP/L2TP-Id/If-Id     : 10.1.1.225 02:00:00:FF:FE:55:55:55
MC-Standby           : No

Session Time Left    : N/A
LCP State            : Opened
IPCP State           : Opened
IPv6CP State         : Opened
PPP MTU              : 1492
PPP Auth-Protocol    : CHAP
PPP User-Name        : sub55@domain1

Subscriber-interface : sub-int-1
Group-interface      : grp-int-1-1

IP Origin            : local-user-db
DNS Origin           : local-user-db
NBNS Origin          : none

Subscriber           : "sub-55"
```

```
Sub-Profile-String   : "sub-prof-3"
SLA-Profile-String   : "sla-prof-5"

--- snipped ---

IP Address           : 10.1.1.225/32
Primary DNS          : 2.2.2.2
Secondary DNS        : N/A
Primary NBNS         : N/A
Secondary NBNS       : N/A
Address-Pool         : N/A

IPv6 Prefix          : N/A
IPv6 Prefix Origin   : none
IPv6 Prefix Pool     : ""
IPv6 Del.Pfx.        : 2001:db8:f101:5500::/56
IPv6 Del.Pfx. Origin : local-user-db
IPv6 Del.Pfx. Pool   : ""
IPv6 Address         : 2001:db8:101:55::55
IPv6 Address Origin  : local-user-db
IPv6 Address Pool    : ""
Primary IPv6 DNS     : 2001:db8:ddd:1::1
Secondary IPv6 DNS   : 2001:db8:ddd:2::1

--- snipped ---

-------------------------------------------------------------------------------
No. of sessions: 1
===============================================================================
*A:BNG-1#
```

The following command shows the lease origin.

```
*A:BNG-1# show service id 1 dhcp6 lease-state session ppp

===============================================================================
DHCP lease state table, service 1
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LeaseTime  Origin   Stdby
-------------------------------------------------------------------------------
2001:db8:101:55::55/128
                00:00:00:55:55:55 1/1/1:111           23h59m17s  UserDb
2001:db8:f101:5500::/56
                00:00:00:55:55:55 1/1/1:111           23h59m17s  UserDb
-------------------------------------------------------------------------------
Number of lease states : 2
===============================================================================
*A:BNG-1#
```

Regular SAPs versus Capture and Managed SAPs

When an LUDB is to be used for regular SAPs, the LUDB must be assigned at the group interface level of a Layer 3 service (IES or VPRN). This LUDB is then used for all SAPs on that group interface, as described in the section Applying an LUDB for ESM.

When an LUDB is to be used for capture and managed SAPs, the LUDB must be assigned at the capture SAPs of the Layer 2 (VPLS) service and at the group interface level of the corresponding Layer 3 service (IES or VPRN).

Because the managed SAPs are dynamically created at the group interface of a Layer 3 service, this service must have its authentication configuration aligned with the Layer 2 service; see Figure 196.

Capture and managed SAPs support IPoE (session and host model) and PPP.

The capture VPLS is defined as follows.

```
configure
    service
        vpls 3 customer 1 create
            stp
                shutdown
            exit
            sap 1/1/2:* capture-sap create
                trigger-packet arp dhcp dhcp6 pppoe rtr-solicit
                dhcp-user-db "ludb-cmsap"
                pppoe-user-db "ludb-cmsap"
                ipoe-session
                    ipoe-session-policy "ipoe-sess-1"
                    user-db "ludb-cmsap"
                    no shutdown
                exit
                msap-defaults
                    group-interface "grp-int-1-1"
                    policy "msap-pol-1"
                    service 2
                exit
            exit
            no shutdown
        exit
```

The VPRN on which the managed SAPs are created is defined as follows.

```
configure
    service
        vprn 2 customer 1 create
            --- snipped ---
            subscriber-interface "sub-int-1" create
                address 10.111.1.254/24
                ipv6
                    delegated-prefix-len 56
```

```
                    subscriber-prefixes
                        prefix 2001:db8:901::/48 wan-host
                        prefix 2001:db8:f901::/48 pd
                    exit
                exit
                group-interface "grp-int-1-1" create
                    --- snipped ---
                    ipoe-session
                        ipoe-session-policy "ipoe-sess-1"
                        sap-session-limit 100
                        user-db "ludb-cmsap"
                        no shutdown
                    exit
                    oper-up-while-empty
                    pppoe
                        session-limit 100
                        user-db "ludb-cmsap"
                        no shutdown
                    exit
                exit
            exit
        exit
```

The msap-defaults needed for creation of the managed SAPs can be taken from the capture SAP, but can also be obtained from an LUDB, as the following example shows. In that case, they overrule the capture SAP msap-defaults.

```
configure
    subscriber-mgmt
        local-user-db "ludb-cmsap" create
            description "LUDB for capture/managed SAPs"
            ipoe
                match-list mac
                host "entry-1" create
                    host-identification
                        mac 00:00:00:01:01:01
                    exit
                    address 10.111.1.101
                    identification-strings 254 create
                        subscriber-id "sub-priv-1"
                        sla-profile-string "sla-prof-3"
                        sub-profile-string "sub-prof-4"
                    exit
                    msap-defaults
                        group-interface "grp-int-1-1"
                        policy "msap-pol-1"
                        service 2
                    exit
                    options
                        subnet-mask 255.255.255.0
                    exit
                    ipv6-address 2001:db8:901:11::11
                    ipv6-delegated-prefix 2001:db8:f901:1100::/56
                    ipv6-delegated-prefix-len 56
                    no shutdown
                exit
            exit
            ppp
```

```
            match-list mac
            host "entry-1" create
                host-identification
                    mac 00:00:00:05:05:05
                exit
                address 10.111.1.105/32
                identification-strings 254 create
                    subscriber-id "sub-05"
                    sla-profile-string "sla-prof-2"
                    sub-profile-string "sub-prof-4"
                exit
                msap-defaults
                    group-interface "grp-int-1-1"
                    policy "msap-pol-1"
                    service 2
                exit
                ipv6-address 2001:db8:901:5::5
                ipv6-delegated-prefix 2001:db8:f901:500::/56
                ipv6-delegated-prefix-len 56
                no shutdown
            exit
        exit
```

Detailed information on managed and capture SAPs is in the Managed SAPs with Routed CO chapter.

The commands to display the subscribers, lease, and session states with the origin codes are the same as in the section PPPoE Authentication, so these are not repeated.

LUDB for ESM as RADIUS Fallback

RADIUS fallback can be triggered in the following situations; see also Figure 195 and Figure 196:

- with the authentication policy directly assigned at the group interface level
- with the authentication policy referenced from an LUDB

For the second case, first-level authentication is performed by the LUDB, and second-level authentication should be performed by the RADIUS server. For both cases, when the RADIUS server is not reachable, fallback happens.

Note: RADIUS fallback is not supported when the LUDB is attached to the group interface or capture SAP via the ipv6 dhcp6 and rtr-solicit contexts.

Although RADIUS fallback applies to both IPoE and PPP, only IPoE is shown in the example that follows.

To demonstrate the use of an LUDB for RADIUS fallback, the configuration of the previous example with capture and managed SAPs is modified, as follows.

```
# the (capture-)VPLS
configure
    service
        vpls 3 customer 1 create
            sap 1/1/2:* capture-sap create
                authentication-policy "auth-pol-1"
                exit
            exit
        exit
    exit

# the VPRN
configure
    service
        vprn 2 customer 1 create
            subscriber-interface "sub-int-1"
                group-interface "grp-int-1-1"
                    authentication-policy "auth-pol-1"
                exit
            exit
        exit
    exit
```

The authentication policy is applied in the VPLS at the SAP level, and in the VPRN at the group interface level. Even with LUDBs assigned in other contexts at that group interface, the authentication policy takes higher priority.

The LUDB used for RADIUS fallback is defined as follows, and both the ipoe and the ppp sections contain a default host entry.

```
configure
    subscriber-mgmt
        local-user-db "ludb-radiusfb" create
            description "LUDB for RADIUS fallback"
            ipoe
                match-list mac
                host "default" create
                    msap-defaults
                        group-interface "grp-int-1-1"
                        policy "msap-pol-1"
                        service 2
                    exit
                    no shutdown
                exit
            exit
            ppp
                match-list username
                    host "default" create
                    msap-defaults
                        group-interface "grp-int-1-1"
```

```
                        policy "msap-pol-1"
                        service 2
                    exit
                    no shutdown
                exit
            exit
            no shutdown
```

The following is the authentication policy from which this LUDB is referenced.

```
configure
    subscriber-mgmt
        authentication-policy "auth-pol-1" create
            fallback-action user-db "ludb-radiusfb"
            radius-server-policy "rsp-1"
        exit
```

The definition of the RADIUS server policy is not relevant so it is not shown.

The following debug configuration applies.

```
debug
    router "Base"
        radius
            packet-type authentication accounting coa
            detail-level high
        exit
    exit
    router "2"
        ip
            dhcp
                detail-level high
                mode egr-ingr-and-dropped
            exit
        exit
    exit
    service
        id 3
            dhcp
                detail-level medium
                mode egr-ingr-and-dropped
            exit
        exit
    exit
    subscriber-mgmt
        local-user-db "ludb-radiusfb"
            detail all
        exit
    exit
exit
```

The following partial debug output shows that when a DHCPv4 user connects, the LUDB "ludb-radiusfb" is accessed after failing to connect to the RADIUS server. Similar debug output appears when connecting through DHCPv6 via IPoE sessions, or PPP.

```
190 2015/10/30 15:55:40.02 CET MINOR: DEBUG #2001 Base SVCMGR
"SVCMGR: RX DHCP Packet
   VPLS 3, SAP 1/1/2:*

   BootRequest to UDP port 67
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:00:00:01:01:01    xid: 0x5

   DHCP options:
   [82] Relay agent information: len = 8
      [1] Circuit-id: 11
      [2] Remote-id: AA
   [53] Message type: Discover
   [255] End
"

191 2015/10/30 15:55:40.03 CET MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  server 192.168.66.66:1812 not reachable"

192 2015/10/30 15:55:40.03 CET MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Access-Request
  user 00:00:00:01:01:01  policy rsp-1
  send failed"

193 2015/10/30 15:55:40.03 CET MINOR: DEBUG #2001 Base LUDB
"LUDB: User lookup success - host found
  mac: 00:00:00:01:01:01

  Host default found in user data base ludb-radiusfb"


194 2015/10/30 15:55:40.03 CET MINOR: DEBUG #2001 vprn2 PIP
"PIP: DHCP
instance 2 (2), interface index 3 (grp-int-1-1),
   received DHCP Boot Request on Interface grp-int-1-1 (1/1/2:123) Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:00:00:01:01:01    xid: 0x5

   DHCP options:
   [82] Relay agent information: len = 8
      [1] Circuit-id: 11
      [2] Remote-id: AA
   [53] Message type: Discover
   [255] End
   Hex Packet Dump:
   --- snipped ---
"

--- snipped ---

201 2015/10/30 15:55:40.07 CET MINOR: DEBUG #2001 vprn2 PIP
"PIP: DHCP
instance 2 (2), interface index 3 (grp-int-1-1),
   transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/2:123) Port 68
```

```
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.111.1.5
   siaddr: 10.111.111.1      giaddr: 10.111.1.254
   chaddr: 00:00:00:01:01:01    xid: 0x5

   DHCP options:
   [82] Relay agent information: len = 8
      [1] Circuit-id: 11
      [2] Remote-id: AA
   [53] Message type: Ack
   [54] DHCP server addr: 10.111.111.1
   [51] Lease time: 864000
   [1] Subnet mask: 255.255.255.0
   [255] End
   --- snipped ---
"
```

In this example, the LUDB accessed on RADIUS fallback defines a default host for ipoe as well as for ppp with msap-defaults only, which means relaying applies where the DHCPv4 and DHCPv6 servers provide the IP addresses and prefixes.

See Figure 193and Figure 194 for the list of supported parameters for IPoE and PPP in the RADIUS fallback scenario.

Operational Considerations and Remarks

The operational considerations listed in the Local User Database Basics chapter still apply.

To maintain backward compatibility with previous software releases, LUDB informational and error messages are sent to the error logs as if they are originating from the DHCPS application (DHCPS #xyz in the preceding outputs).

Conclusion

LUDBs offer a self-contained method of providing ESM data locally stored on the router, so that no external database is needed for supporting authentication. In case authentication relies on an AAA/RADIUS server that fails, an LUDB can provide the ESM data instead through RADIUS fallback. LUDBs can be used on regular, managed, and capture SAPs.

Managed SAPs with Routed CO

This chapter provides information about Managed SAPs with Routed CO.

Topics in this chapter include:

Applicability

This chapter is applicable to the 7750 SR7/12 with IOM2 or higher (for BRAS functionality) with Chassis mode B or higher (for Routed Central Office (CO) model) and the 7710/7750 SR-c12 and was tested on release 12.0.R1. Routed CO is supported on 7450 ESS-7 or ESS-12 in mixed-mode since 8.0.R1. The 7750 SR-c4 is supported from 8.0.R4 and higher.

This note is related only to the use of IPv4.

MSAPs are also supported with Bridged CO model and on the 7450, however, applicable configuration information is beyond the scope of this document.

Overview

Managed Service Access Point (MSAP) allows the use of policies and a SAP template for the creation of a SAP. As part of the MSAP feature, individual SAPs are created along with the subscriber host with minimal configuration on the BRAS node. Creation of a managed SAP is triggered by a DHCP-DISCOVER and/or a PPPoE-PADI message. In this case, the authentication response message not only returns the subscriber host attributes, but also the managed SAP policy and service ID.These latter two parameters are used by the system to create the subscriber SAP with default settings as indicated in the managed SAP policy and then assigning it to the corresponding VPN service. In this model, each subscriber is defined with its own VLAN. This feature uses authentication mechanisms supported by the node to create a SAP.

When enabled, receiving a triggering packet initiates 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 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.

Knowledge of TPSDA (Triple Play Service Delivery Architecture) and functionality is assumed throughout this document.

The network topology is displayed in Figure 198. The configuration consists of one 7750 SR-12 acting as a BNG with BRAS functionality.

Figure 198 **Network Topology**

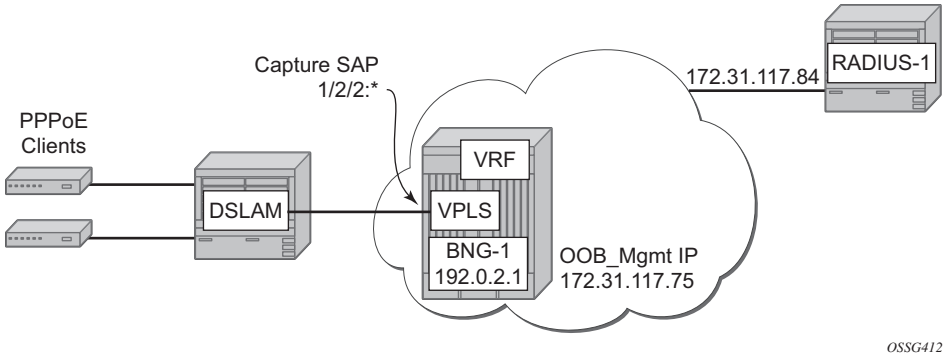

Capture SAP

A capture SAP is used to capture triggering packets and initiate RADIUS authentication. This SAP is defined in a similar way to a default SAP but does not forward traffic.

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 encapsulation. Managed SAPs and regular SAPs can co-exist on the same port and in the same service.

The capture SAP is used if a more specific match for the Q or Q-in-Q tags is not found by the traffic classification on 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.

The following are examples for supported modes:

SAP 1/2/2:*	for dot1Q
SAP 1/2/2:Q1.*	for QinQ (where Q1 > 0)

The MSAP created will have a single tag (for dot1q) or both q-tags (for qinq) that arrived in the original packet if authenticated by RADIUS.

While MSAPs are supported in both routed CO and bridged CO Triple Play Service Delivery Architecture (TPSDA) models, the triggering SAP must be created in a VPLS service.

Triggering Packets

DHCP discover (or requests if re-authentication is configured) for DHCP clients. The managed SAP lifetime is defined by the lease time.

PPPoE PADI for the PPPoE client. The MSAP lifetime is defined by the session time. The MSAP is installed after the IP address is provided.

ARP packets as trigger packets within a capture SAP. ARP trigger packets can be used for static IP hosts. The managed SAP lifetime is defined by the ARP entry lifetime and is subject to the same ARP entry refresh mechanisms as other ARP entries.

All trigger types can be combined on a SAP supporting DHCP, PPPoE and ARP hosts. In this chapter, a PPPoE client is used.

RADIUS Authentication and Vendor Specifc Attributes (VSAs) for MSAP

An MSAP is created in the service-id context that is returned from RADIUS. The RADIUS attribute Alc-MSAP-Serv-Id refers to the service in which the MSAP is created.

In a Routed CO scenario, the MSAP is created in a group-interface context. The group-interface name is returned from RADIUS attribute Alc-MSAP-Interface and must exist in the provided service for the MSAP to be installed.

The MSAP parameters are defined in the creation policy. The policy name is returned from RADIUS in the attribute Alc-MSAP-Policy in order for the MSAP to be created.

Configuration

Configure RADIUS Authentication Policy “authentication-1”

The following output shows a RADIUS authentication policy configuration defining “authentication-1”.

```
configure subscriber-mgmt
     authentication-policy "authentication-1" create
         radius-authentication-server
             source-address 172.31.117.75
             router "management"
             server 1 address 172.31.117.84 secret ALU
         exit
         pppoe-access-method pap-chap
         include-radius-attribute
             remote-id
             nas-identifier
             mac-address
         exit
     exit
exit
```

Where, management routing instance and the out-of-band and IP address 172.31.117.75 are used as a source address to communicate authentication messages between the BNG and the RADIUS server. The RADIUS server IP address is 172.31.117.84. Up to five servers can be configured. When having multiple servers two possible access algorithms can be configured to access the list of RADIUS servers, **direct** or **round-robin**.

The value of secret is ALU which is case sensitive and must be configured on Clients.conf file on the RADIUS server in advance. Up to 20 characters in length are possible.

The authentication method used in our example is PAP/CHAP, so the pap-chap value is used for the pppoe-access-method.

The user’s remote-id and mac-address are sent as well the nas-identifier into the access request message towards the RADIUS.

By default, the RADIUS authentication messages are send over port 1812 but can be overridden by adding an explicit port setting to the **server** command.

```
configure subscriber-mgmt
     authentication-policy "authentication-1" create
         radius-authentication-server
             server 1 address 172.31.117.84 secret ALU port <value>
```

Configure a RADIUS Accounting Policy

This example configures radius-accounting-policy "accounting-1".

```
configure subscriber-mgmt
    radius-accounting-policy "accounting-1" create
        update-interval 10
        include-radius-attribute
            framed-ip-addr
            subscriber-id
            circuit-id
            remote-id
            nas-port-id
            nas-identifier
            sub-profile
            sla-profile
            user-name
            no detailed-acct-attributes
            std-acct-attributes
        exit
        session-id-format number
        radius-accounting-server
            router "management"
            server 1 address 172.31.117.84 secret ALU
        exit
    exit
exit
```

Where, accounting updates are sent every 10 mins (the default update-interval is 5 minutes). The accounting session-id-format in this example is a number (40 HEX character string).

```
SESSION ID [44] 40 00000001024100000000064000000034B090B2D
```

Whereas, session-id-format <description> can be used in this case. The session-id-format is as follows:

```
<subscriber>@<sapid>@<SLA-profile>_<creation-time>

SESSION ID [44] 50 user1@1/2/2:100@sla-profile-2M_2009/11/22 11:56:25
```

Since std-acct-attributes is used, only the total number of octets/packets in ingress and egress directions are sent.

ALU VSAs are used for accounting, in such case, detailed accounting values for each queue (in case multiple queues for the subscriber can be used) and the in-profile and the out-profile values are shown. This feature can be enabled by adding **no std-acct-attribute**, which is the default.

By default, the RADIUS accounting messages are sent over port 1813 but can be overridden by adding an explicit port setting in addition to the **server** command.

```
configure subscriber-mgmt
     radius-accounting-policy accounting-1 create
         radius-authentication-server
             server 1 address 172.31.117.84 secret ALU port <value>
```

Configure an QoS SAP Ingress Policy

Configure QoS SAP ingress policy where shaping and SAP egress policy performs shaping and remarking. Values for dot1p and dscp are used as examples.

```
configure qos
     sap-ingress 20 create
         description "64K_upstream"
         queue 1 create
             rate 64
         exit
     exit
     sap-ingress 30 create
         description "128K_upstream"
         queue 1 create
             rate 128
         exit
     exit
     sap-ingress 40 create
         description "256K_upstream"
         queue 1 create
             rate 256
         exit
     exit
     sap-ingress 50 create
         description "512K_upstream"
         queue 1 create
             rate 512
```

```
            exit
        exit
        sap-egress 20 create
            description "256K_downstream"
            queue 1 create
                rate 256
            exit
            fc be create
                queue 1
                dot1p 5
                dscp ef
            exit
        exit
        sap-egress 30 create
            description "512K_downstream"
            queue 1 create
                rate 512
            exit
            fc be create
                queue 1
                dot1p 4
                dscp af21
            exit
        exit
        sap-egress 40 create
            description "1M_downstream"
            queue 1 create
                rate 1024
            exit
            fc be create
                queue 1
                dot1p 5
                dscp ef
            exit
        exit
        sap-egress 50 create
            description "2M_downstream"
            queue 1 create
                rate 2048
            exit
            fc be create
                queue 1
                dot1p 3
                dscp cs1
            exit
        exit
    exit
```

Configure Enhanced Subscriber Management Parameters

Four SLA profiles are configured where the downstream speed is four times the upstream speed and the SLA profile will be named with the downstream speed.

Also, a subscriber profile will be configured to initiate RADIUS accounting and doing SLA profile mapping.

```
configure subscriber-mgmt
        sla-profile "sla-profile-1M" create
            ingress
                qos 40 shared-queuing
                exit
            exit
            egress
                qos 40
                exit
                no qos-marking-from-sap
            exit
        exit
        sla-profile "sla-profile-256K" create
            ingress
                qos 20 shared-queuing
                exit
            exit
            egress
                qos 20
                exit
                no qos-marking-from-sap
            exit
        exit
        sla-profile "sla-profile-2M" create
            ingress
                qos 50 shared-queuing
                exit
            exit
            egress
                qos 50
                exit
                no qos-marking-from-sap
            exit
        exit
        sla-profile "sla-profile-512K" create
            ingress
                qos 30 shared-queuing
                exit
            exit
            egress
                qos 30
                exit
                no qos-marking-from-sap
            exit
        exit
        sub-profile "sub-profile-default" create
            radius-accounting-policy "accounting-1"
            sla-profile-map
                use-direct-map-as-default
            exit
        exit
        sub-ident-policy "sub-id-default" create
            sub-profile-map
                use-direct-map-as-default
            exit
```

```
            sla-profile-map
                use-direct-map-as-default
            exit
        exit
exit
```

Configure an MSAP Policy

MSAP policies contain the configuration template (parameters) to be used for MSAP creation and the necessary information to complete the subscriber identification process.

The MSAP policy that will be used is either returned by RADIUS in the access-accept message during authentication phase if this MSAP policy is already configured under subscriber management context, or else the default MSAP policy will be used instead.

```
configure subscriber-mgmt
        msap-policy "msap-ISP1" create
            sub-sla-mgmt
                def-sub-id use-sap-id
                def-sub-profile "sub-profile-default"
                def-sla-profile "sla-profile-512K"
                sub-ident-policy "sub-id-default"
                single-sub-parameters
                     profiled-traffic-only
                exit
            exit
        exit
        msap-policy "msap-default" create
            sub-sla-mgmt
                def-sub-id use-sap-id
                def-sub-profile "sub-profile-default"
                def-sla-profile "sla-profile-256K"
                sub-ident-policy "sub-id-default"
                single-sub-parameters
                     profiled-traffic-only
                exit
            exit
        exit
exit
```

If managed routes are required for a certain subscriber, add the following command under msap-policy. The default anti-spoof is **ip-mac**. Managed routes are out of the scope of this document.

```
configure subscriber-mgmt
        msap-policy "msap-ISP1" create
            ies-vprn-only-sap-parameters
                anti-spoof nh-mac
            exit
        exit
```

Configure a VPLS Service with a Capture SAP

Configure a VPLS service with capture SAP and define the triggering packet types. The **trigger-packet** and **authentication-policy** commands are mandatory within the capture SAP. Additionally, the **cpu-protection** command can be added to enable CPU protection policies

```
configure
    service
        vpls 1 customer 1 create
            description "VPLS for Capture SAPs"
            stp
                shutdown
            exit
            sap 1/2/2:* capture-sap create
                description "capture SAP for MSAP creation on port 1/2/2"
                trigger-packet arp dhcp pppoe
                msap-defaults
                    policy "msap-default"
                exit
                authentication-policy "authentication-1"
            exit
            no shutdown
        exit
```

Verify the details of capture SAP:

```
A:BNG# show service id 1 sap 1/2/2:* detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 1
SAP                : 1/2/2:*                   Encap             : q-tag
Description        : capture SAP for MSAP creation on port 1/2/2
Admin State        : Up                        Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 03/20/2014 11:28:26
Last Mgmt Change   : 03/20/2014 11:28:09
Sub Type           : capture
Triggers           : arp dhcp pppoe
Dot1Q Ethertype    : 0x8100                    QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

<snipped>

Egr MCast Grp      :
Auth Policy        : authentication-1
DHCP User Db       : None
PPP Policy         : None
PPP User Db        : None
PPPoE Policy       : default
PPPoE User Db      : None
DHCPv6 User Db     : None
```

```
<snipped>
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                 Octets
CPM Ingress           : 0                       0

Forwarding Engine Stats
Dropped               : 0                       0

DHCP Capture Stats
Received              : 0
Redirected            : 0
Dropped               : 0

<snipped>

PPP Capture Stats
Received              : 0
Redirected            : 0
Dropped               : 0

Rtr-Sol Capture Stats
Received              : 0
Redirected            : 0
Dropped               : 0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets
No entries found
===============================================================================
A:BNG#
```

Note: The dropped packets are those that are non triggering packets. Also, there are no SAP queues instantiated for a capture SAP.

Configuration Scenario — Routed CO/VLAN-Per-Subscriber (PPPOE)

The following output shows a Routed CO configuration example.

```
configure service vprn 2
        route-distinguisher 65000:2
        subscriber-interface "sub-int-1" create
            address 10.255.255.254/8
            group-interface "group-int-1" create
```

```
                    description "ROUTED CO MSAP VLAN X"
                    authentication-policy "authentication-1"
                    pppoe
                        session-limit 2000
                        no shutdown
                    exit
                exit
            exit
            no shutdown
    exit
```

Note: The number of PPPoE sessions can be controlled under a group interface by applying the **pppoe session-limit** command.

Initially, since no MSAPs are present, the operational state of both the subscriber interface and group interface context are down.

```
*A:BNG# show router 2 interface
===============================================================================
Interface Table (Service: 2)
===============================================================================
Interface-Name                   Adm         Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                    PfxState
-------------------------------------------------------------------------------
group-int-1                      Up          Down/Down   VPRN G* n/a
sub-int-1                        Up          Down/Down   VPRN S* subscriber
   10.255.255.254/8                                              n/a
-------------------------------------------------------------------------------
Interfaces : 2
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:BNG#
```

To allow the subscriber interface to consider this group interface to be operationally enabled without any active SAPs, the following command can be added to the configuration (this would be useful in order to propagate the subnet interface address into a routing protocol) :

```
configure service vprn 2
        subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    oper-up-while-empty

*A:BNG# show router 2 interface
===============================================================================
Interface Table (Service: 2)
===============================================================================
Interface-Name                   Adm         Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                    PfxState
-------------------------------------------------------------------------------
group-int-1                      Up          Down/Down   VPRN G* n/a
```

```
sub-int-1                          Up          Up/Down     VPRN S* subscriber
   10.255.255.254/8                                                n/a
-------------------------------------------------------------------------------
Interfaces : 2
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:BNG#
```

Note the status of the group interface once the first MSAP is created.

Configure RADIUS User Files

The following entry is an example of a user entry in a RADIUS users file for FreeRadius server.

```
"user1@ISP1.com"      Cleartext-Password := "user1_pass"
                      Alc-Subsc-ID-Str := "%{ADSL-Agent-Remote-Id}",
                      Alc-SLA-Prof-Str == "sla-profile-2M",
                      Alc-MSAP-Serv-ID = 2,
                      Alc-MSAP-Policy == "msap-ISP1",
                      Alc-MSAP-Interface == "group-int-1",
                      Framed-IP-Address = 10.255.0.1,
                      Alc-Primary-DNS = 67.138.54.100,
                      Alc-Secondary-DNS = 207.225.209.66
```

So when the PPPoE user sends the correct username and password, the RADIUS will accept the access message and returns the correct VPRN service id 2, the correct group interface group-int-1 the MSAP policy to use msap-ISP1.

In case there are no MSAP policy returned from RADIUS, the default MSAP policy sap-default under the capture SAP will be used instead.

In the above entry, the PPPoE user will have its IP address and DNS assigned by RADIUS as well. The DNS values are examples for public Free DNSs.

Connect PPPoE “user1”

Connect PPPoE user1, initiate a PPPoE session on VLAN 1 and verify PPPoE session establishment.

```
*A:BNG# show service id 2 pppoe session
===============================================================================
PPPoE sessions for svc-id 2
===============================================================================
```

```
Sap Id              Mac Address       Sid    Up Time         Type
    IP/L2TP-Id/Interface-Id                                    MC-Stdby
-------------------------------------------------------------------------------
[1/2/2:1]           00:00:86:1c:79:a1 1      0d 00:00:29     local
    10.255.0.1
-------------------------------------------------------------------------------
Number of sessions   : 1
===============================================================================
*A:BNG#
```

The PPPoE session is established successfully and the user obtained the IP and subscriber strings from the RADIUS server.

In order to differentiate between the MSAP and the normal SAP, the MSAP will be shown between square brackets [1/2/2:1] in the show commands

Verify Subscriber Values

Verify subscriber values returned from RADIUS for user1.

```
*A:BNG# show service id 2 pppoe session ip-address 10.255.0.1 detail
===============================================================================
PPPoE sessions for svc-id 2
===============================================================================
Sap Id              Mac Address       Sid    Up Time         Type
    IP/L2TP-Id/Interface-Id                                    MC-Stdby
-------------------------------------------------------------------------------
[1/2/2:1]           00:00:86:1c:79:a1 1      0d 00:00:42     local
    10.255.0.1

LCP State            : Opened
IPCP State           : Opened
IPv6CP State         : Initial
PPP MTU              : 1492
PPP Auth-Protocol    : CHAP
PPP User-Name        : user1@ISP1.com

Subscriber-interface : sub-int-1
Group-interface      : group-int-1

IP Origin            : radius
DNS Origin           : radius
NBNS Origin          : none

Subscriber           : "user1"
Sub-Profile-String   : ""
SLA-Profile-String   : "sla-profile-2M"
ANCP-String          : ""
Int-Dest-Id          : ""
App-Profile-String   : ""
Category-Map-Name    : ""
Acct-Session-Id      : "EA4BFF00000000532AD1CD"
Sap-Session-Index    : 1
```

```
IP Address           : 10.255.0.1/32
Primary DNS          : 67.138.54.100
Secondary DNS        : 207.225.209.66
Primary NBNS         : N/A
Secondary NBNS       : N/A
Address-Pool         : N/A

IPv6 Prefix          : N/A
IPv6 Prefix Origin   : none
IPv6 Prefix Pool     : ""
IPv6 Del.Pfx.        : N/A
IPv6 Del.Pfx. Origin : none
IPv6 Del.Pfx. Pool   : ""
IPv6 Address         : N/A
IPv6 Address Origin  : none
IPv6 Address Pool    : ""
Primary IPv6 DNS     : N/A
Secondary IPv6 DNS   : N/A

Circuit-Id           : DSLAM1_1/1/1/1:0.35
Remote-Id            : user1

Radius Session-TO    : N/A
Radius Class         :
Radius User-Name     : user1@ISP1.com
Logical-Line-Id      :
Service-Name         :
-------------------------------------------------------------------------------
Number of sessions   : 1
===============================================================================
*A:BNG#
```

Check the Actual Values

Check the actual values used by user1, subscriber profile, SLA profile, VPRN and group interface association, the subscriber queues statistics and others.

```
*A:BNG# show service active-subscribers subscriber "user1" detail
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber user1 (sub-profile-default)
-------------------------------------------------------------------------------
I. Sched. Policy : N/A
E. Sched. Policy : N/A                          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.  : accounting-1
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 : "user1"
Sub. Int Dest Id : ""
Igmp Rate Adj    : N/A
RADIUS Rate-Limit: N/A
Oper-Rate-Limit  : Maximum
* indicates that the corresponding row element may have been truncated.
-------------------------------------------------------------------------------
(1) SLA Profile Instance
    - sap:[1/2/2:1] (VPRN 2 - group-int-1)
    - sla:sla-profile-2M
-------------------------------------------------------------------------------
Description          : (Not Specified)
Host Limit           : No Limit
Egr Sched-Policy     : N/A
Ingress Qos-Policy   : 50                    Egress Qos-Policy : 50
Ingress Queuing Type : Shared-queuing (Not Applicable to Policer)
Ingr IP Fltr-Id      : N/A                   Egr IP Fltr-Id    : N/A
Ingr IPv6 Fltr-Id    : N/A                   Egr IPv6 Fltr-Id  : N/A
Ingress Report-Rate  : Maximum
Egress Report-Rate   : Maximum
Egress Remarking     : from SLA Profile Qos
Credit Control Pol.  : N/A
Category Map         : (Not Specified)
Use ing L2TP DSCP    : false
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.255.0.1
                00:00:86:1c:79:a1 1         IPCP

-------------------------------------------------------------------------
SLA Profile Instance statistics
-------------------------------------------------------------------------
                        Packets                 Octets

Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Off. Uncolor          : 0                       0
Off. Managed          : 0                       0

Queueing Stats (Ingress QoS Policy 50)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Queueing Stats (Egress QoS Policy 50)
Dro. InProf           : 0                       0
```

```
Dro. OutProf          : 0                      0
For. InProf           : 0                      0
For. OutProf          : 1                      64

-------------------------------------------------------------------------------
SLA Profile Instance per Queue statistics
-------------------------------------------------------------------------------
                        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
Dro. InProf           : 0                      0
Dro. OutProf          : 0                      0
For. InProf           : 0                      0
For. OutProf          : 1                      64

-------------------------------------------------------------------------------
*A:BNG#
```

Where, the subscriber id is **user1**, subscriber profile is sub-profile-default.

Note: The RADIUS did not return subscriber profile string, so the system will use the **def-sub-profile** configured under the msap-policy msap-ISP1.

Another command can also be used to show less detail in a hierarchical form.

```
*A:BNG# show service active-subscribers hierarchy subscriber "user1"
===============================================================================
Active Subscriber hierarchy
===============================================================================
-- user1 (sub-profile-default)
   |
   |-- sap:[1/2/2:1] - sla:sla-profile-2M
   |   |
   |   |-- 10.255.0.1
   |   |   00:00:86:1c:79:a1 - 1 (IPCP)
   |   |

===============================================================================
*A:BNG#
```

Verify that the IPv4 state of the group interface is now up.

```
*A:BNG# show router 2 interface
===============================================================================
Interface Table (Service: 2)
===============================================================================
Interface-Name                   Adm         Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                    PfxState
-------------------------------------------------------------------------------
group-int-1                      Up          Up/Down     VPRN G* 1/2/2
sub-int-1                        Up          Up/Down     VPRN S* subscriber
   10.255.255.254/8                                              n/a
-------------------------------------------------------------------------------
Interfaces : 2
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:BNG#
```

Verify the capture service id (VPLS), capture SAP and the msap policy used to created user1 and the SAP sub type.

```
*A:BNG# show service id 2 sap 1/2/2:1 detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 2
SAP                : 1/2/2:1                  Encap             : q-tag
Description        : Managed SAP - Capture Svc 1 1/2/2:*
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 03/20/2014 11:28:08
Last Mgmt Change   : 03/20/2014 11:32:29
Sub Type           : managed
Capture Service Id : 1                        Capture SAP       : 1/2/2:*
MSAP Policy        : msap-ISP1
Dot1Q Ethertype    : 0x8100                   QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

<snip>
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                         Packets                 Octets
No entries found
===============================================================================
*A:BNG#
```

The sub type shows managed for MSAPs, whereas regular for normal saps (a SAP created manually under a group-interface).

MSAP with Redundant Configurations

MSAPs are High Availability (HA) enabled (there is no service impact following a CPM failover). In addition, the MSAPs are also stored in the subscriber management persistence file (if enabled), allowing the MSAPs to be recreated after a reboot.

MSAPs can be used in dual-homed BNG scenarios with multi-chassis LAG, multi-chassis ring and subscriber router redundancy protocol.

MSAP QoS Notes

An MSAP is always created with default QoS policies.

```
*A:BNG# show service id 2 sap 1/2/2:1 detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 2
SAP                : 1/2/2:1                  Encap             : q-tag
Description        : Managed SAP - Capture Svc 1 1/2/2:*
Admin State        : Up                       Oper State        : Up

<snip>
-------------------------------------------------------------------------------
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      : default                  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)
-------------------------------------------------------------------------------
Subscriber Management
-------------------------------------------------------------------------------
Admin State        : Up                       MAC DA Hashing    : False
Def Sub-Id         : Use sap-id (1/2/2:1)
Def Sub-Profile    : sub-profile-default
Def SLA-Profile    : sla-profile-512K

<snip>
===============================================================================
*A:BNG#
```

QoS Egress Remarking

In order to have remarking for egress traffic for MSAP taken from SLA profile, use **no qos-marking-from-sap** command.

```
configure subscriber-mgmt
        ...
        sla-profile "sla-profile-512K" create
            ingress
                qos 30 shared-queuing
                exit
            exit
            egress
                qos 30
                exit
                no qos-marking-from-sap
            exit
        exit
```

By default, the egress QoS marking for subscriber-host traffic is derived from the SAP-egress QoS policy associated with the corresponding SAP rather than the SLA profile associated with the corresponding subscriber-host. As a consequence, no egress QoS marking (for example, dot1p marking was set to 0, DSCP/PREC field is unchanged) is performed for traffic transmitted on an MSAP because per default, SAP-egress policy one (1) was attached to every MSAP.

Queue Optimization

Shared queuing can be used to optimize queues on ingress direction.

```
configure subscriber-mgmt
        ...
        sla-profile "sla-profile-512K" create
            ingress
                qos 30 shared-queuing
                exit
            exit
```

The SAP queues will not be instantiated when using the following option in the msap-policy.

```
configure subscriber-mgmt
    msap-policy "msap-ISP1" create
        sub-sla-mgmt
```

```
            single-sub-parameters
                profiled-traffic-only
            exit
        exit
    exit
```

Configuration Tips

The authentication policy used in the capture SAP must be the same as the policy used for the managed SAP.

The managed SAP will not be created if the group-interface name returned from RADIUS points to a different authentication policy other than the policy defined by the capture SAP.

```
configure
    service
        vpls 1
            --- snip ---
            sap 1/2/2:* capture-sap create
                --- snip ---
                authentication-policy "authentication-1"
            exit
            no shutdown
        exit


configure
    service
        vprn 2
            subscriber-interface "sub-int-1"create
                --- snip ---
                group-interface "group-int-1" create
                    authentication-policy "authentication-2"
                    --- snip ---
                exit
            exit
            no shutdown
        exit
```

This can be seen in log 99:

```
84 2014/03/20 11:35:37.80 UTC WARNING: PPPOE #2001 Base PPPoE session failure
"PPPoE session failure on SAP 1/2/2:* in service 1 -
 [00:00:86:1c:79:a1,1,user1@ISP1.com] MSAP group-interface "group-int-
1" RADIUS auth-policy "authentication-2" differs from capture SAP"

83 2014/03/20 11:35:37.80 UTC MINOR: SVCMGR #2214 Base Managed SAP creation failure
"The system could not create Managed SAP:1/2/
2:1, MAC:00:00:86:1c:79:a1, Capturing SAP:1/2/
2:*, Service:1. Description: MSAP group-interface "group-int-1" RADIUS auth-
policy "authentication-2" differs from capture SAP"
```

On the 7750 SR, enable debug for PPPoE and RADIUS packets to help in case there is a problem in session establishment:

```
debug
    router "management"
        radius
            packet-type authentication accounting coa
            detail-level medium
        exit
    exit
    service
        id 1
            ppp
                packet
                    mode egr-ingr-and-dropped
                    detail-level medium
                    discovery
                    ppp
                exit
            exit
        exit
        id 2
            ppp
                packet
                    mode egr-ingr-and-dropped
                    detail-level medium
                    discovery
                    ppp
                    dhcp-client
                exit
            exit
        exit
    exit


configure
    log
        log-id 1
            from debug-trace
            to session
        exit
    exit
exit
```

Disconnect/connect user1 then check the RADIUS access request/accept and accounting messages from the debug output.

```
14 2014/03/20 12:38:42.04 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.31.117.84:1812 id 26 len 184 vrid 4095 pol authenticatio
n-1
    USER NAME [1] 14 user1@ISP1.com
    NAS IP ADDRESS [4] 4 172.31.117.75
    SERVICE TYPE [6] 4 Framed(2)
    FRAMED PROTOCOL [7] 4 PPP(1)
    CHAP PASSWORD [3] 17 1 0xb54dcb79d5de3fd6cff4ad7b98ac3598
```

```
    CHAP CHALLENGE [60] 51 0xa52131167c5ff2adef841422767b7acb458de8c95c2bf2c7185
8fe09a1794f471a80dd975f50c44fd4d8f0cb54ea9719f781e2
    VSA [26] 7 DSL(3561)
      AGENT REMOTE ID [2] 5 user1
    NAS PORT TYPE [61] 4 PPPoEoVLAN(33)
    NAS PORT ID [87] 7 1/2/2:1
    NAS IDENTIFIER [32] 3 BNG
    VSA [26] 19 Alcatel(6527)
      CHADDR [27] 17 00:00:86:1c:79:a1
"

15 2014/03/20 12:38:42.04 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Receive
  Access-Accept(2) id 26 len 133 from 172.31.117.84:1812 vrid 4095 pol authentic
ation-1
    VSA [26] 7 Alcatel(6527)
      SUBSC ID STR [11] 5 user1
    VSA [26] 16 Alcatel(6527)
      SLA PROF STR [13] 14 sla-profile-2M
    VSA [26] 6 Alcatel(6527)
      MSAP SERVICE ID [31] 4 2
    VSA [26] 11 Alcatel(6527)
      MSAP POLICY [32] 9 msap-ISP1
    VSA [26] 13 Alcatel(6527)
      MSAP INTERFACE [33] 11 group-int-1
    FRAMED IP ADDRESS [8] 4 10.255.0.1
    VSA [26] 6 Alcatel(6527)
      PRIMARY DNS [9] 4 67.138.54.100
    VSA [26] 6 Alcatel(6527)
      SECONDARY DNS [10] 4 207.225.209.66
" "
```

The 7750 sends also accounting request message to the RADIUS accounting server.

```
23 2014/03/20 12:38:42.11 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Transmit
  Accounting-Request(4) 172.31.117.84:1813 id 15 len 200 vrid 4095 pol accountin
g-1
    STATUS TYPE [40] 4 Start(1)
    NAS IP ADDRESS [4] 4 172.31.117.75
    USER NAME [1] 14 user1@ISP1.com
    SERVICE TYPE [6] 4 Framed(2)
    FRAMED PROTOCOL [7] 4 PPP(1)
    FRAMED IP ADDRESS [8] 4 10.255.0.1
    NAS IDENTIFIER [32] 3 BNG
    SESSION ID [44] 22 EA4BFF0000000E532AE152
    EVENT TIMESTAMP [55] 4 1395319122
    NAS PORT TYPE [61] 4 PPPoEoVLAN(33)
    NAS PORT ID [87] 7 1/2/2:1
    VSA [26] 28 DSL(3561)
      AGENT CIRCUIT ID [1] 19 DSLAM1_1/1/1/1:0.35
      AGENT REMOTE ID [2] 5 user1
    VSA [26] 44 Alcatel(6527)
      SUBSC ID STR [11] 5 user1
      SUBSC PROF STR [12] 19 sub-profile-default
      SLA PROF STR [13] 14 sla-profile-2M
"
```

After 10 mins (update interval) the 7750 sends accounting Interim updates with the same session ID including the counter values for total input and output octets/ packets for user1.

```
25 2014/03/20 12:48:47.65 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Transmit
  Accounting-Request(4) 172.31.117.84:1813 id 16 len 230 vrid 4095 pol accountin
g-1
    STATUS TYPE [40] 4 Interim-Update(3)
    NAS IP ADDRESS [4] 4 172.31.117.75
    USER NAME [1] 14 user1@ISP1.com
    SERVICE TYPE [6] 4 Framed(2)
    FRAMED PROTOCOL [7] 4 PPP(1)
    FRAMED IP ADDRESS [8] 4 10.255.0.1
    NAS IDENTIFIER [32] 3 BNG
    SESSION ID [44] 22 EA4BFF0000000E532AE152
    SESSION TIME [46] 4 606
    EVENT TIMESTAMP [55] 4 1395319727
    NAS PORT TYPE [61] 4 PPPoEoVLAN(33)
    NAS PORT ID [87] 7 1/2/2:1
    VSA [26] 28 DSL(3561)
      AGENT CIRCUIT ID [1] 19 DSLAM1_1/1/1/1:0.35
      AGENT REMOTE ID [2] 5 user1
    VSA [26] 44 Alcatel(6527)
      SUBSC ID STR [11] 5 user1
      SUBSC PROF STR [12] 19 sub-profile-default
      SLA PROF STR [13] 14 sla-profile-2M
    INPUT PACKETS [47] 4 0
    INPUT OCTETS [42] 4 0
    OUTPUT PACKETS [48] 4 11
    OUTPUT OCTETS [43] 4 704
"

26 2014/03/20 12:48:47.65 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Receive
  Accounting-Response(5) id 16 len 20 from 172.31.117.84:1813 vrid 4095 pol acco
unting-1
" "
```

To verify the MSAP policies and associations of MSAPs created, use the following commands:

```
*A:BNG# show subscriber-mgmt msap-policy
===============================================================================
Managed SAP Policies
===============================================================================
Name                            Num    Description
                                MSAPs
-------------------------------------------------------------------------------
msap-ISP1                       1      (Not Specified)
msap-default                    0      (Not Specified)
-------------------------------------------------------------------------------
Number of MSAP Policies : 2
Number of MSAPs         : 1
===============================================================================
*A:BNG#
```

```
*A:BNG# show subscriber-mgmt msap-policy "msap-ISP1" association
===============================================================================
MSAP Policy Associations
===============================================================================
Service-Id : 2 (VPRN)
 - SAP : [1/2/2:1]
-------------------------------------------------------------------------------
Number of associated MSAPs: 1
===============================================================================
*A:BNG#
```

To check all MSAPs created and associations to services.

```
*A:BNG# show service sap-using msap
===============================================================================
Service Access Points
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
[1/2/2:1]                       2          1     none    1     none   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
Number of Managed SAPs : 1, indicated by [<sap-id>]
-------------------------------------------------------------------------------
===============================================================================
*A:BNG#
```

It is possible to use a **tools** command to update an existing MSAP when a specific msap-policy has changed.

```
A:BNG-1# tools perform subscriber-mgmt eval-msap ?
  - eval-msap { policy <msap-policy-name< | msap <sap-id> }

<msap-policy-name>    :  [32 chars max]
<sap-id>              :  <port-id|lag-id>:qtag1
                         <port-id|lag-id>:qtag1.qtag2
```

To delete an MSAP.

```
A:BNG-1# clear service id 2 msap 1/2/2:1

166 2014/03/20 11:48:21.39 UTC INDETERMINATE: LOGGER #2010 Base Clear SVCMGR
"Clear function clearSvcIdMsap has been run with parameters: svc-id="2" sap-id="1/2/
2:1".  The completion result is: success.  Additional error text, if any, is: "
```

To delete all MSAPs associated with a certain MSAP policy use the following command:

```
A:BNG-1# clear service id 2 msap-policy msap-ISP1

168 2014/03/20 11:48:32.15 UTC INDETERMINATE: LOGGER #2010 Base Clear SVCMGR
"Clear function clearSvcIdMsapPlcy has been run with parameters: svc-id="2" policy-
name="msap-
ISP1".  The completion result is: success.  Additional error text, if any, is: "
```

Conclusion

MSAP allows dynamic creation of SAPs which results in:

- Less provisioning.
- Less possibility for introducing provisioning errors.
- Reduced configuration file.

Multi-Chassis Ring Layer 2 with Enhanced Subscriber Management

This chapter provides information about MC-Ring Layer 2 with Enhanced Subscriber Management (ESM).

Topics in this chapter include:

Applicability

These configuration notes are applicable to all of the 7450, 7750 and 7710 SR series and was tested on Release 7.0R5. The 7750 SR-c4 is supported from 8.0R4 and higher.

MC-Ring L2 with Enhanced Subscriber Management (ESM) was introduced in SR OS 6.0. There are no other pre-requisites for the 7450, 7750 and 7710 SR this configuration.

Summary

Multi-Chassis Ring (MC-ring) is an extension for dual homing support in TPSDA (Triple Play Service Delivery Architecture) networks based on Layer 2 CO (Layer 2 Central Office) model. The extension addresses networks where multiple access nodes (for example, DSLAMs, GPON OLT) are connected in a single ring.

MC Ring Layer 2 ESM is considered an extension or evolution for ring topologies of the Multi-Chassis LAG dual-homing solution used for directly connected access nodes.

MC Ring Layer 2 CO is documented in the Triple Play Enhanced Subscriber Management / Dual Homing section of the Triple Play Guide.

Overview

MC-Ring

Figure 199 shows a typical ring-based configuration of network model based on the Layer 2 CO model.

Individual rings of access nodes are aggregated at the Broadband Service Aggregator (BSA) level in one (or multiple) Virtual Private LAN Service (VPLS) service(s). At higher aggregation levels, Broadband Service Router (BSR) individual BSAs are connected to a Layer 3 interface (Internet Enhanced Service (IES) or Virtual Private Routed Network (VPRN)) by means of a spoke SDP (Service Distribution Point) termination. Every Layer 3 interface at the BSR level aggregates all subscribers in one or more subnets. ESM is performed in the BSAs.

Figure 199 **MC-Ring Layer 2 CO Dual Homing**

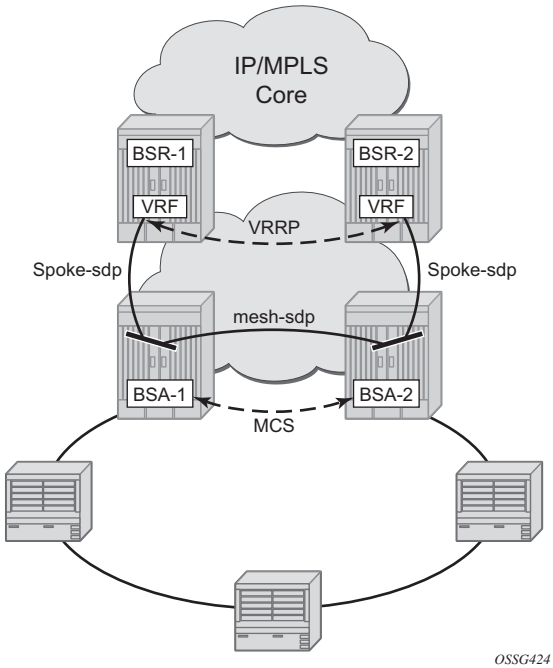

The key functional components of the MC-Ring Layer 2 CO dual-homing redundancy solution are:

1. Mirroring of the subscriber state between the two BSAs performing subscriber management using the multi-chassis synchronization (MCS) protocol (BSA-1 and BSA-2 in Figure 199).
2. Ring control between the two BSAs, using the following mechanisms:
   - In-band bi-directional forwarding detection (BFD) between the BSAs over the ring to monitor ring integrity and detect failures. A BFD session between BSA-1 and BSA-2 runs through the access ring using a dedicated IES/VPRN interface configured on both BSAs. This connection uses a separate VLAN throughout the ring (access nodes provide transparent bridging for this VLAN).
   - Out-of-band communication between the BSA nodes to exchange information about the reachability of individual access nodes, and to verify the configuration consistency of the ring. The information on configuration is synchronized through MCS (this use of MCS is related to, but independent to, MCS for subscriber-state synchronization mentioned above).
3. Ring Node Connectivity Verification (RNCV). Each BSA uses RNCV to detect the reachability of individual Access Nodes. It is used after a ring failure to determine which BSA should handle the traffic of each Access Node. RNCV uses ARP requests to “ping” individual ANs, which must be configured with an IP address for this purpose.
4. Per-subscriber attribute (intermediate destination ID) for the BSA to know the Access Node each subscriber is connected to (assigned through RADIUS or DHCP/Python).
5. VRRP in the BSRs to provide a redundant default gateway to the CPEs/Home Gateways. BSRs do not perform any subscriber management functions.

The operation of dual homing at the BSA level will be covered based on two underlying mechanisms.

MC-Ring Layer 2 CO Operation

To describe the functional behavior and operation of the dual-homing concept in a ring, it is best to sub-divide the description into the following three parts:

- Steady-state operation with ring fully closed
- Transition to ring-broken state
- Transition from ring-broken to steady state operation

Steady-State Operation of Dual-Homed Ring

Figure 200 illustrates the detailed operation of the dual-homed ring. The steady-state is achieved under following conditions:

- Both nodes are configured in a consistent way
- The MCS peering relation is up
- In-Band Ring Control Connection (IB-RCC) is in an operationally UP state.

Note: This connection is set-up using BFD session between IP interfaces on BSA-1 and BSA-2

Figure 200 **Dual homing Under Steady-State Condition**

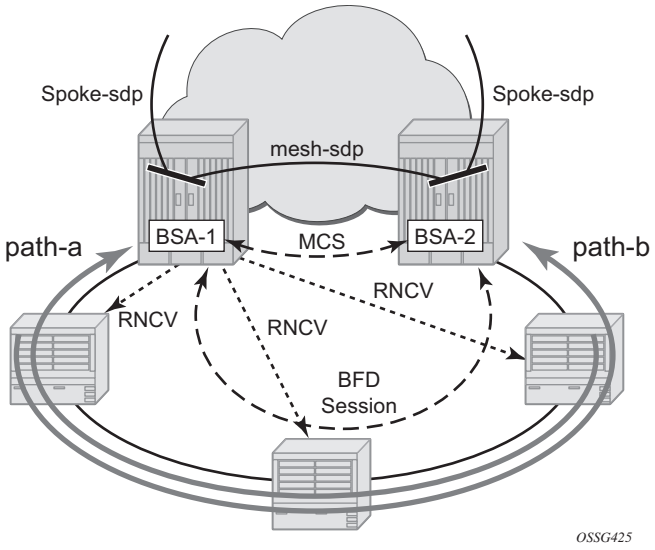

Under the above conditions, the ring is fully closed and every access node (the **ring-node**) has two possible paths toward the VPLS core. Figure 200 refers to them as **path-a** and **path-b**. In order to avoid the loop created by the ring, only one of the paths may be used by any given ring-node for any given VLAN. The assignment of the individual VLANs to path-a or path-b, respectively, has to be provisioned on both BSAs in a consistent manner. The BSA with a lower IP address in the interface used for BFD will be the master for the VLANs associated with path-a and standby for the VLANs associated with path-b.

The following summarizes the behavior of the master and standby BSA for a given path (a or b) and the VLANs configured for that path:

Master BSA for the VLANs/SAPs associated with a given path (a or b):

- SAPs associated with the path are operationally up and FDB entries for sub-hosts point to the corresponding SAP
- Master BSA for a path performs RNCV checks to all ring nodes. RNCV failures trigger an alarm but do **not** trigger a switchover
- The ARP reply agent replies to ARP requests for subscriber-hosts associated with ring-nodes for which the BSA is master

Standby BSA:

- All SAPs associated with the path for which the BSA is standby will be operationally down and all FDB entries of subscriber hosts associated with those SAPs will point towards an SDP connecting to the master BSA.

Broken-Ring Operation and the Transition to this State

Figure 201 illustrates the scenario with the broken ring (link failure or ring-node failure). This state is reached under following conditions:

- Both nodes are configured in a consistent way
- The MCS peering relation is up
- IB-RCC is in operationally down state

In this situation, every ring node has only one access path towards the VPLS core and hence the path-a and path-b notion has no real meaning in this situation.

From a functional point of view, both BSAs are now the master for the ring-nodes they can reach. For all hosts behind unreachable ring-nodes, the corresponding subscriber host FDB (Forwarding DataBase) entries will point to SDP pointing to the other head-end ring PE.

Figure 201 **Broken Ring State**

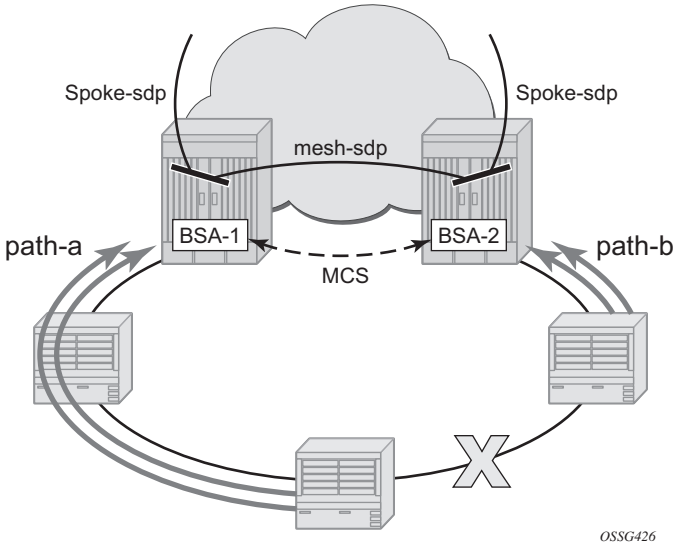

Transition from Broken to Closed Ring State

MC-ring operates in a revertive mode. This means that whenever the ring connectivity is restored, the BSA with the lower IP address in the IB-RCC communication channel will become master of path-a and slave for ring-b with vice versa operations on the other BSA.

The actions in this case are straightforward. After restoration of BFD session, the master functionality is assumed by respective BSAs. The FDB tables are updated according to the master/standby role of the given BSA and FDB population messages will be sent accordingly.

Configuration

Figure 202 **Network Topology**

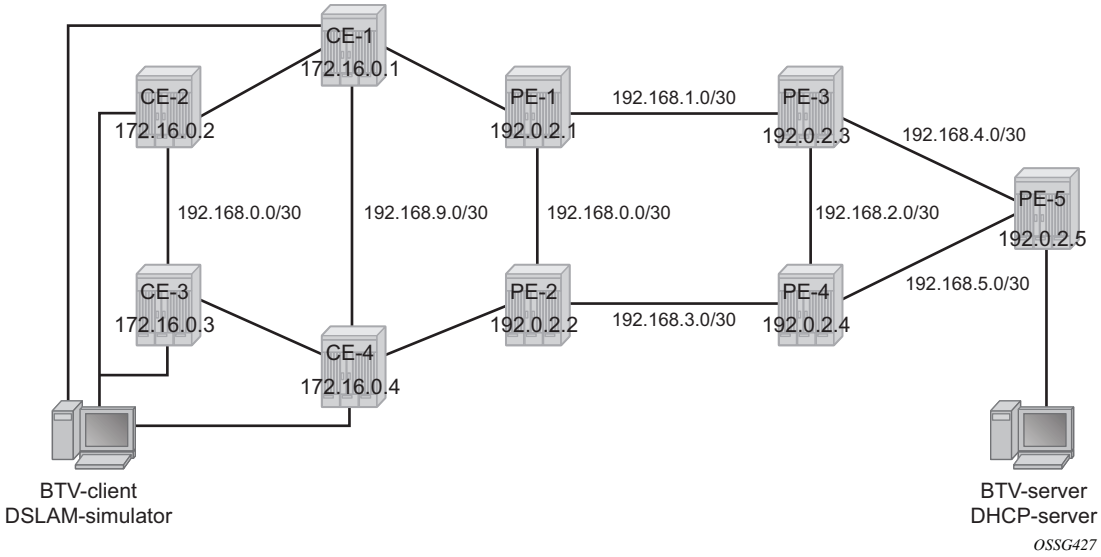

The network topology is displayed in Figure 202. The setup consists of two BSA nodes (PE-1 and PE-2), two BSR nodes (PE-3, PE-4) and another PE router (PE-5). A setup with one BSR node and two BSA nodes can also be used, but in this example, two (2) BSR nodes were used to show the typical Layer 2-CO setup with VRRP and Protocol Independent Multicast (PIM) on the BSRs. The access ring consists of four CE nodes.

Base Topology

This section assumes that the following base configuration has been implemented on the PE:

- Cards, MDAs and ports configured
- Interfaces configured
- IGP (Interior Gateway Protocol) configured and converged LDP (Label Distribution Protocol) configured on the interfaces between PE-3-PE-1, PE-1-PE-2 and PE-2-PE-4
- T-LDP (Targeted-LDP) configured on PE-1, PE-2, PE-3, PE-4
- SDPs configured between PE-3-PE-1, PE-1-PE-2 and PE-2-PE-4

Note: You can choose between OSPF and IS-IS as the IGP. Both LDP and RSVP (Resource Reservation Protocol) can be used for signaling the transport MPLS labels. Alternatively, GRE (Generic Routing Encapsulation) can be used for the transport tunnels. In this example, OSPF is configured and LDP SDPs are used.

7750 SR routers are used as ring-nodes to simulate Access Nodes. Ring-nodes can be any L2 device that support VLANs and have the ability to connect an IP interface to one VLAN (required for RNCV).

NTP Configuration

Time must be the same on the redundant NSAs (PE-1 and PE-2); otherwise, the lease times will be different on both nodes. NTP can be used:

```
configure system time
            ntp
                server 10.30.30.30 prefer
                no shutdown
            exit
exit all
```

MC-Ring Configuration

Configure a VPLS service on the CE routers for the In-Band Ring Control connection (BFD packets between PE-1 and PE-2 traversing the ring). This VPLS service will also be used for RNCV.

Note: An Epipe service can also be used in case the service is only used for BFD. In that case, a separate service is required for the RNCV.

In this example, QinQ encapsulation is used and BFD and RNCV packets will use VLAN tag 1. Note that dot1q encapsulation can also be used.

The configuration of CE-1 is shown below. A similar configuration is required on the other CE routers.

```
configure
    port 1/1/1
        ethernet
            mode access
            encap-type qinq
```

```
        exit
        no shutdown
    exit
    port 1/1/2
        ethernet
            mode access
            encap-type qinq
        exit
        no shutdown
    exit
    service
        vpls 1 customer 1 create
            interface "lo1" create
                address 172.16.0.1/24
            exit
            sap 1/1/1:1.0 create
            exit
            sap 1/1/2:1.0 create
            exit
            no shutdown
        exit
exit all
```

An interface **lo1** is created in the VPLS. This interface will be used for RNCV.

On the BSA nodes (PE-1 and PE-2), configure an IES services that will originate BFD and RNCV messages. On PE-1:

```
configure
    port 1/1/1
        ethernet
            mode access
            encap-type qinq
        exit
        no shutdown
    exit
    service
        ies 1 customer 1 create
            interface "bfd-rncv1" create
                address 172.16.0.251/24
                bfd 100 receive 100 multiplier 3
                sap 1/1/1:1.0 create
                exit
            exit
            no shutdown
        exit
exit all
```

On PE-2:

```
configure
    port 1/1/2
        ethernet
            mode access
            encap-type qinq
        exit
```

```
        no shutdown
    exit
    service
        ies 1 customer 1 create
            interface "bfd-rncv1" create
                address 172.16.0.252/24
                bfd 100 receive 100 multiplier 3
                sap 1/1/2:1.0 create
                exit
            exit
            no shutdown
        exit
exit all
```

In-Band Ring Control Connection BFD is always originated on an IES or VPRN service on the BSA nodes. RNCV messages can be originated from the same IES/ VPRN service or from a separate service.

Configure Multi-Chassis Synchronization (MCS) on the BSA nodes. Enable MCS for **igmp-snooping**, **mc-ring** and **sub-mgmt**. The configuration of PE-1 is shown below. The configuration of PE-2 is similar.

```
configure redundancy multi-chassis peer 192.0.2.2 create
                sync
                    igmp-snooping
                    mc-ring
                    sub-mgmt
                    port 1/1/1 sync-tag "l2ring1" create
                    exit
                    no shutdown
                exit
                no shutdown
exit all
```

Add the MC-ring configuration on the BSA nodes and link the Ring Control Connection BFD session to the IES service created before.

```
configure redundancy multi-chassis peer 192.0.2.2 create
                mc-ring
                    ring "l2ring1" create
                        in-band-control-path
                            service-id 1
                            interface "bfd-rncv1"
                            dst-ip 172.16.0.252
                        exit
                        no shutdown
                    exit
                exit
                no shutdown
exit all
```

Note that the ring name is exactly the same as the sync-tag already configured.

At this moment, the MC-ring should be up:

```
A:PE-1# show redundancy multi-chassis mc-ring peer 192.0.2.2 ring l2ring1 detail
===============================================================================
Multi-Chassis MC-Ring Detailed Information
===============================================================================
Peer            : 192.0.2.2
Ring Type       : Layer 2
Sync Tag        : l2ring1
Port ID         : 1/1/1
Admin State     : inService
Oper State      : connected
Admin Change    : 11/05/2009 21:17:54
Oper Change     : 11/05/2009 21:17:54
Failure Reason  : None
Control B Path  : No
-------------------------------------------------------------------------------
In Band Control Path
-------------------------------------------------------------------------------
Service ID      : 1
Interface Name  : bfd-rncv1
Oper State      : connected
Dest IP         : 172.16.0.252
Src  IP         : 172.16.0.251
...
```

Next, configure under MCS the ring nodes on PE-1 and PE-2. This configuration will be used for the RNCV. The ring node configuration for CE-1 is shown below:

```
configure redundancy multi-chassis peer 192.0.2.2 mc-ring ring l2ring1
                        ring-node "CE-1" create
                            connectivity-verify
                                dst-ip 172.16.0.1
                                interval 1
                                service-id 1
                                vlan 1
                                no shutdown
                            exit
                        exit
exit all
```

Note: The **interval** parameter is the interval used to check the reachability of the CE nodes (configurable from 1 to 6000 minutes). A BFD failure will also be a trigger for a reachability check.

The configuration on PE-2 is identical and the configuration of the other ring nodes is similar.

Configure VLAN 3 to take path-b (default is path-a).

```
configure redundancy multi-chassis peer 192.0.2.2 mc-ring ring l2ring1
                        path-b
                            range 3-3
                        exit
exit all
```

MC-Ring Verification

Verify that MCS is up and running:

```
A:PE-1# show redundancy multi-chassis all
===============================================================================
Multi-Chassis Peers
===============================================================================
Peer IP         Src IP          Auth       Peer Admin   MC-Ring Oper MC-EP Adm
 MCS Admin       MCS Oper        MCS State  MC-LAG Adm   MC-LAG Oper
-------------------------------------------------------------------------------
192.0.2.2       192.0.2.1       None       Enabled      inService    --
 Enabled         Enabled         inSync     Disabled     Disabled
===============================================================================
```

The output shows that MCS is administrative and operational up and that both peers are synchronized. MC-ring is operational in service.

The following output shows more detailed information about the configured ring:

```
A:PE-1# show redundancy multi-chassis mc-ring peer 192.0.2.2 ring l2ring1 detail
===============================================================================
Multi-Chassis MC-Ring Detailed Information
===============================================================================
Peer           : 192.0.2.2
Ring Type      : Layer 2
Sync Tag       : l2ring1
Port ID        : 1/1/1
Admin State    : inService
Oper State     : connected
Admin Change   : 11/05/2009 21:22:29
Oper Change    : 11/05/2009 21:22:29
Failure Reason : None
Control B Path : No
-------------------------------------------------------------------------------
In Band Control Path
-------------------------------------------------------------------------------
Service ID     : 1
Interface Name : bfd-rncv1
Oper State     : connected
Dest IP        : 172.16.0.252
Src  IP        : 172.16.0.251
Debounce State : inService
Debounce Max   : 10 s
Debounce Guard : 60 s
```

```
-------------------------------------------------------------------------------
VLAN Managed Range
-------------------------------------------------------------------------------
full range
-------------------------------------------------------------------------------
VLAN Map B Path Provisioned
-------------------------------------------------------------------------------
range 3-3
-------------------------------------------------------------------------------
VLAN Map Excluded Path Provisioned
-------------------------------------------------------------------------------
no range
-------------------------------------------------------------------------------
VLAN Map B Path Operational
-------------------------------------------------------------------------------
range 3-3
-------------------------------------------------------------------------------
VLAN Map Excluded Path Operational
-------------------------------------------------------------------------------
no range
```

The output above shows that the ring **l2ring1** on port 1/1/1 is in service and connected. Ring-node Connectivity Check (BFD) is running on **interface bfd-rncv1** in service **1**. All VLANs on port 1/1/1 use path-a (default) except VLAN 3, which uses path-b.

The BSA peer with the lower IP address (the master) will put the SAPs configured for path-a in operational up state while the SAPs configured in path-b will be put in operational down state. The other BSA peer will do the reverse. This is done to prevent loops in the ring.

Note: No VLANs are configured to be excluded from the paths. If a VLAN is configured to be excluded from the paths, the respective SAPs of this VLAN on both BSA nodes will be operationally up.

Check which ring-nodes are configured and connected:

```
A:PE-1# show redundancy multi-chassis mc-ring peer 192.0.2.2 ring l2ring1 ring-node
===============================================================================
MC-Ring Node entries
===============================================================================
Name                            Loc Oper St.      Failure Reason
  In Use                          Rem Oper St.
-------------------------------------------------------------------------------
CE-1                            connected         None
  No                              notTested
CE-2                            connected         None
  No                              notTested
CE-3                            connected         None
  No                              notTested
CE-4                            connected         None
  No                              notTested
-------------------------------------------------------------------------------
```

```
No. of MC-Ring Node entries: 4
===============================================================================
A:PE-1#
```

The output above shows that four ring-nodes are connected. Only the master will send RNCV messages to the ring-nodes. As soon as a ring failure occurs, the BFD session goes down and both BSA nodes send out RNCV messages to see which ring-nodes are connected.

Note: When the reachability of the CE nodes is determined, the RNCV messages will be sent at a continuous interval (see above).

More detail about each ring-node can be provided with following command:

```
A:PE-1# show redundancy multi-chassis mc-ring peer 192.0.2.2 ring l2ring1 ring-
node CE-1 detail
===============================================================================
Multi-Chassis MC-Ring Node Detailed Information
===============================================================================
Peer            : 192.0.2.2
Sync Tag        : l2ring1
Node Name       : CE-1
Oper State Loc  : connected
Oper State Rem  : notTested
In Use          : False
Admin Change    : 11/05/2009 21:21:30
Oper Change     : 11/05/2009 21:22:32
Failure Reason  : None
-------------------------------------------------------------------------------
Ring Node Connectivity Verification
-------------------------------------------------------------------------------
Admin State     : inService
Service ID      : 1
Encap Value     : 1
Dest IP         : 172.16.0.1
Src  IP         : None
Interval        : 1 minutes
Src MAC         : None
===============================================================================
A:PE-1#
```

Unicast Services Configuration

Figure 203 shows the logical setup of the services that will be created. A mesh SDP is used between PE-1 and PE-2. A spoke SDP could also be used.

Figure 203 **Unicast Services — Logical Topology**

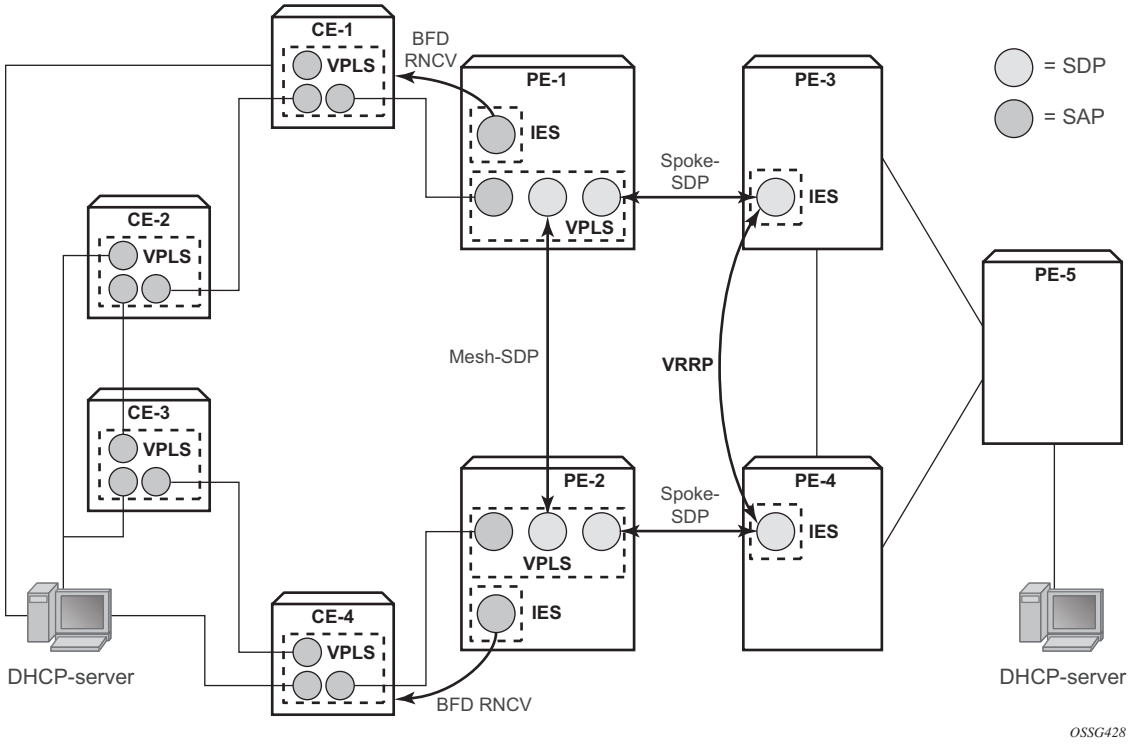

In this setup, two unicast services (VPLS 2 and VPLS 3) are created. VPLS 2 uses path-a and VPLS 3 uses path-b.

The services use a mesh SDP between PE-1 and PE-2 and a spoke-SDP between PE-1/PE-3 and another spoke SDP between PE-2/PE-4. On PE-3 and PE-4 a spoke SDP terminated IES is configured with a VRRP default gateway address. The VRRP packets are switched through the BSAs.

IGMP snooping and ESM are configured on both services. ESM is required since the BSA node must know which ring-node each subscriber is connected to. In this setup, the intermediate destination identifier (int_dest_id) will be returned by Option 254 in x Dynamic Host Control Protocol (DHCP). ESM configuration details are outside the scope of this document. Refer to the appropriate platform OS Triple Play Guide.

The following output shows the configuration of VPLS 2 on PE-1:

```
configure service
       vpls 2 customer 1 create
            description "VLAN_2"
            split-horizon-group "RSHG" residential-group create
            exit
            sap 1/1/1:2.* split-horizon-group "RSHG" create
                dhcp
                    snoop
```

```
                    lease-populate 10
                    no shutdown
                exit
                anti-spoof ip-mac
                sub-sla-mgmt
                    def-sub-profile "initial"
                    def-sla-profile "initial"
                    sub-ident-policy "speedy"
                    multi-sub-sap 10
                    no shutdown
                exit
            exit
            mesh-sdp 12:2 create
                dhcp
                    snoop
                exit
            exit
            spoke-sdp 13:2 create
                dhcp
                    snoop
                exit
            exit
            no shutdown
        exit
exit all
```

The configuration of VPLS 3 is similar. The configuration of VPLS 2 and 3 are similar on PE-2.

The output below shows the subscriber management configuration on PE-1. Similar configuration is required on PE-2.

```
configure subscriber-mgmt
        sla-profile "initial" create
        exit
        sub-profile "initial" create
        exit
        sub-ident-policy "speedy" create
            strings-from-option 254
        exit
exit all
```

Also, configure VLAN 2 and 3 on all the ring-nodes. This configuration is straightforward. Below is an example of the VLAN 2 configuration on CE-1. In this example, the client is connect to port 1/1/3 and should send packets with an outer tag of 2:

```
configure service
        vpls 2 customer 1 create
            sap 1/1/1:2.* create
            exit
            sap 1/1/2:2.* create
            exit
            sap 1/1/3:2.* create
            exit
```

```
            no shutdown
        exit
exit all
```

In the example, VLAN 2.* and 3.* is used to allow for transparent transport of the customer VLAN.

The IES service where VPLS 2 terminates looks like this on BSR PE-3:

```
configure service
        ies 2 customer 1 create
            interface "VLAN_2" create
                address 10.0.2.3/24
                dhcp
                    server 10.10.10.10
                    trusted
                    no shutdown
                exit
                ip-mtu 1500
                vrrp 2
                    backup 10.0.2.254
                    ping-reply
                exit
                local-proxy-arp
                spoke-sdp 31:2 create
                exit
            exit
            no shutdown
        exit
exit all
```

And on PE-4:

```
configure service
        ies 2 customer 1 create
            interface "VLAN_2" create
                address 10.0.2.4/24
                dhcp
                    server 10.10.10.10
                    trusted
                    no shutdown
                exit
                ip-mtu 1500
                vrrp 2
                    backup 10.0.2.254
                    ping-reply
                exit
                local-proxy-arp
                spoke-sdp 42:2 create
                exit
            exit
            no shutdown
        exit
exit all
```

Notice that the ip-mtu must be set to match the vc-mtu signaled by the other side of the spoke-SDP. Otherwise, the service will be operationally down with a ServiceMTUMismatch.

Note also in the configuration that DHCP relay is done by configuring a DHCP server under the IES interface.

The configuration of IES 3 on PE-3 and PE-4 is similar..

On PE-5 an IES interface is configured to the DHCP-server. The interface is configured with a Dot1Q encapsulated port because this port will be also be used for an interface to the multicast server.

```
configure service
        ies 2 customer 1 create
            interface "dhcp-server" create
                address 192.168.6.1/30
                sap 1/1/3:2 create
                exit
            exit
            no shutdown
        exit
exit all
```

Unicast Services Verification

Request an IP address on VLAN 2 on CE-1. On BSA routers PE-1 and PE-2 following DHCP info can be checked:

```
A:PE-1# show service id 2 dhcp lease-state
===============================================================================
DHCP lease state table, service 2
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LifeTime   Origin   Stdby
-------------------------------------------------------------------------------
10.0.2.107      00:00:64:01:01:02 1/1/1:2.*           09d07h38m  DHCP
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
```

ESM show commands can be used to obtain the subscriber identity, IP address, MAC address and SLA-profile:

```
A:PE-1# show service active-subscribers
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber subscriber_1_vlan_2 (initial)
-------------------------------------------------------------------------------
```

```
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:1/1/1:2.* - sla:initial
-------------------------------------------------------------------------------
IP Address      MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.0.2.107      00:00:64:01:01:02 N/A       DHCP

-------------------------------------------------------------------------------
Number of active subscribers : 1
===============================================================================
```

The following command gives more information about a specific subscriber:

```
A:PE-1# show service active-subscribers subscriber subscriber_1_vlan_2 detail

===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber subscriber_1_vlan_2 (initial)
-------------------------------------------------------------------------------
I. Sched. Policy : N/A
E. Sched. Policy : N/A                              E. Agg Rate Limit: Max
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
Sub. ANCP-String : "subscriber_1_vlan_2"
Sub. Int Dest Id : "CE-1"
Host Trk Rate Adj: N/A
-------------------------------------------------------------------------------
(1) SLA Profile Instance
    - sap:1/1/1:2.* (VPLS 2)
    - sla:initial
-------------------------------------------------------------------------------
Description          : (Not Specified)
Host Limit           : No Limit
Ingress Qos-Policy   : 1                      Egress Qos-Policy : 1
Ingress Queuing Type : Service-queuing
Ingress Filter-Id    : N/A                    Egress Filter-Id  : N/A
Ingress Report-Rate  : N/A
Egress Report-Rate   : N/A
Egress Remarking     : from Sap Qos
Credit Control Pol.  : N/A
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
IP Address      MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.0.2.107      00:00:64:01:01:02 N/A       DHCP
```

The output above gives more details about which ring-node the customer is connected to (Sub. Int Dest Id : CE-1), which QoS policies are applied, statistics of each queue,

MCS Verification

Check if the two redundant peers are in sync and check the detailed MCS info:

```
A:PE-1# show redundancy multi-chassis sync peer 192.0.2.2
===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 192.0.2.2
Description             : (Not Specified)
Authentication          : Disabled
Source IP Address       : 192.0.2.1
Admin State             : Enabled
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     : IGMPSnooping SUBMGMT RING
Sync Admin State        : Up
Sync Oper State         : Up
DB Sync State           : inSync
Num Entries             : 11
Lcl Deleted Entries     : 0
Alarm Entries           : 0
Rem Num Entries         : 11
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             : 1
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 1
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
-------------------------------------------------------------------------------
Application             : mcRing
Num Entries             : 10
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 10
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
-------------------------------------------------------------------------------
Application             : mldSnooping
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
-------------------------------------------------------------------------------
Application             : dhcpServer
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
-------------------------------------------------------------------------------
Application             : subHostTrk
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
-------------------------------------------------------------------------------
===============================================================================
A:PE-1#
```

The output shows that both MCS peers are in sync and that entries are populated for MC-ring and Subscriber Management (DHCP lease states).

Notice that the lease states are also populated on PE-2:

```
A:PE-2# show service id 2 dhcp lease-state
===============================================================================
DHCP lease state table, service 2
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
```

```
                                                        LifeTime   Origin   Stdby
-------------------------------------------------------------------------------
10.0.2.107      00:00:64:01:01:02 1/1/2:2.*             09d06h16m  DHCP     Yes
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
A:PE-2#
```

The output is similar to the output on PE-1 except that on PE-2 the flag MC-Stdby is set to yes, which implies that this node is in standby mode for this VLAN.

This can be verified by looking at the status of the SAP on PE-1 and PE-2. On PE-1 the SAP is operationally up:

```
A:PE-1# show service id 2 sap 1/1/1:2.*
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 2
SAP                : 1/1/1:2.*                Encap             : qinq
QinQ Dot1p         : Default
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 11/06/2009 17:17:26
Last Mgmt Change   : 11/04/2009 22:51:15
===============================================================================
```

On PE-2 the situation is different:

```
A:PE-2# show service id 2 sap 1/1/2:2.*
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 2
SAP                : 1/1/2:2.*                Encap             : qinq
QinQ Dot1p         : Default
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Down
Flags              : StandByForMcRing
Multi Svc Site     : None
Last Status Change : 11/06/2009 18:07:30
Last Mgmt Change   : 11/04/2009 23:43:16
===============================================================================
```

Notice that the SAP on PE-2 is operationally down and that a flag is set: StandByForMcRing.

The situation is reversed for VLAN 3 since it was configured for path-b; here the SAP on PE-1 is operationally down and the SAP on PE-2 is operationally up.

Ring Failure Verification

In case a ring failure occurs (either ring link failure or ring-node failure), the IB-RCC BFD session between PE-1 and PE-2 will go down and both nodes will put the SAP in operational up state.

Break the link between CE-2 and CE-1.

Observe that the ring is now in a **broken** state:

```
A:PE-1# show redundancy multi-chassis mc-ring peer 192.0.2.2
===============================================================================
MC-Ring entries
===============================================================================
Sync Tag                        Oper State      Failure Reason
-------------------------------------------------------------------------------
l2ring1                         broken          None
-------------------------------------------------------------------------------
No. of MC-Ring entries: 1
===============================================================================
```

The following command shows the BSA to ring nodes connections:

```
A:PE-1# show redundancy multi-chassis mc-ring peer 192.0.2.2 ring l2ring1 ring-node
===============================================================================
MC-Ring Node entries
===============================================================================
Name                            Loc Oper St.      Failure Reason
  In Use                          Rem Oper St.
-------------------------------------------------------------------------------
CE-1                            connected         None
  Yes                             disconnected
CE-2                            disconnected      None
  No                              connected
CE-3                            disconnected      None
  No                              connected
CE-4                            disconnected      None
  No                              connected
-------------------------------------------------------------------------------
No. of MC-Ring Node entries: 4
===============================================================================
```

The output shows that CE-1 is connected to PE-1. CE-2, CE-3 and CE4 are connected to PE-2.

Notice that the SAPs on both PE-1 and PE-2 are now operationally up. This can be checked with the **show service id 2 sap 1/1/1:2.*** command on PE1 and the **show service id 2 sap 1/1/2:2.*** command on PE2.

Following the ring failure, the BSA nodes send a MAC address withdrawal message to all the SDP peers configured in the VPLS.

Multicast Service Configuration

In general, BTV (Broadcast TV) services are delivered to the DSLAM on a different VLAN using a different VPLS service. The DSLAM can send/relay IGMP group membership messages if it wants to receive a multicast stream. At the BSA level (PE-1 and PE-2), IGMP snooping is configured. At the BSR level (PE-3, PE-4) and the PE where the video source is located, PIM is configured and IGMP is configured on the IES service facing the BSA ring. The BSA ring consists of concatenated spoke SDPs. The spoke SDP ring is not closed between PE-3 and PE-4 to avoid a loop. Figure 204 shows the logical setup for the multicast service with a MC-ring.

Figure 204 **Multicast Service — Logical Setup**

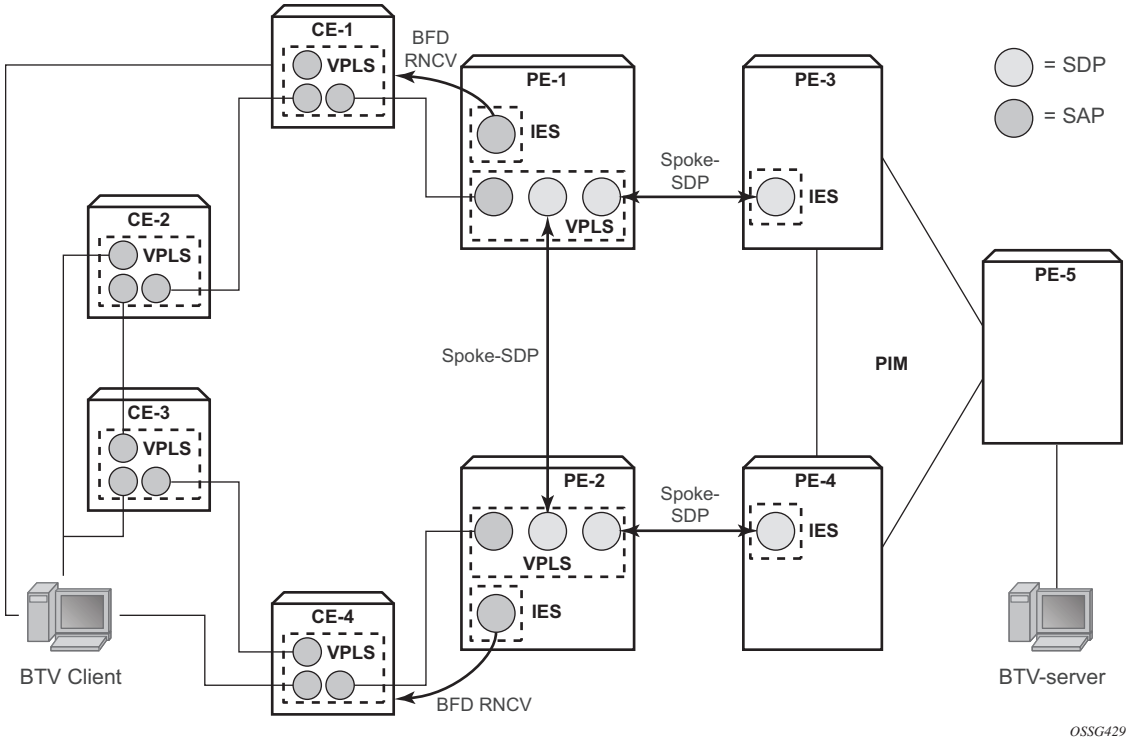

Configure an interface to the multicast server on PE-5:

```
configure service
        ies 3 customer 1 create
            interface "mcast-server" create
                address 192.168.7.1/30
                sap 1/1/3:3 create
                exit
            exit
            no shutdown
        exit
exit all
```

Configure a VPLS service on PE-1 and PE-2. The configuration on PE-1 is shown below. The VPLS configuration on PE-2 is similar.

```
configure service
        vpls 4 customer 1 create
            description "Multicast VLAN"
            igmp-snooping
                no shutdown
            exit
            sap 1/1/1:4.* create
            exit
            spoke-sdp 12:4 create
                igmp-snooping
                    mrouter-port
                exit
            exit
            spoke-sdp 13:4 create
                igmp-snooping
                    mrouter-port
                exit
            exit
            no shutdown
        exit
exit all
```

Notice that igmp-snooping has been enabled and that the spoke SDPs are configured as mrouter-ports in order to forward the IGMP joins to the BSRs.

Configure a spoke SDP terminated IES service on PE-3:

```
configure service
        ies 4 customer 1 create
            interface "btv-dst" create
                address 10.0.4.3/24
                ip-mtu 1500
                spoke-sdp 31:4 create
                exit
            exit
            no shutdown
        exit
exit all
```

On PE-4:

```
configure service
        ies 4 customer 1 create
            interface "btv-dst" create
                address 10.0.4.4/24
                ip-mtu 1500
                spoke-sdp 42:4 create
                exit
            exit
            no shutdown
        exit
exit all
```

Notice that also here the **ip-mtu** must be configured to bring the service up. The **ip-mtu** is required to have an MTU match on the spoke SDP between the IES service and the VPLS service.

Configure PIM on PE-3, PE-4 and PE-5. The configuration on PE-3 is shown below. The PIM configuration on PE-4 and PE-5 is similar.

```
configure router pim
            interface "system"
            exit
            interface "int-PE-3-PE-4"
                priority 10
            exit
            interface "int-PE-3-PE-5"
            exit
            interface "btv-dst"
            exit
            rp
                static
                    address 192.0.2.5
                        group-prefix 224.0.0.0/4
                    exit
                exit
            exit
exit all
```

Note: PE-5 is statically configured as the RP. This is just an example. Different configurations can be used.

Configure IGMP on PE-3/PE-4:

```
configure router igmp interface btv-dst no shutdown
```

The service should also be configured on all ring nodes. Below, the configuration on CE-2 is shown. Similar configurations are required on the other ring-nodes.

```
configure service
        vpls 4 customer 1 create
            sap 1/1/1:4.* create
            exit
            sap 1/1/2:4.* create
            exit
            sap 1/1/3:4.* create
            exit
            no shutdown
        exit
exit all
```

Multicast Service Verification

Configure the multicast server to send one or more multicast streams. Have the BTV client connected to CE-2 send an IGMP join message for this multicast stream.

On the BSA routers (PE-1/PE-2), IGMP snooping can be checked:

```
A:PE-1# show service id 4 mfib
===============================================================================
Multicast FIB, Service 4
===============================================================================
Source Address  Group Address         Sap/Sdp Id               Svc Id   Fwd/Blk
-------------------------------------------------------------------------------
*               *                     sdp:12:4                 Local    Fwd
                                      sdp:13:4                 Local    Fwd
*               225.1.1.1             sap:1/1/1:4.*            Local    Fwd
                                      sdp:12:4                 Local    Fwd
                                      sdp:13:4                 Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
```

On PE-3/PE-4 the IGMP groups can be checked:

```
A:PE-3# show router igmp group
===============================================================================
IGMP Groups
===============================================================================

(*,225.1.1.1)                          Up Time : 0d 00:01:03
    Fwd List  : btv-dst
-------------------------------------------------------------------------------
(*,G)/(S,G) Entries : 1
===============================================================================
```

Following command shows that the MCS peers are synchronized and that there is one IGMP entry on both peers:

```
A:PE-1# show redundancy multi-chassis sync peer 192.0.2.2 detail
===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 192.0.2.2
Description             : (Not Specified)
Authentication          : Disabled
Source IP Address       : 192.0.2.1
Admin State             : Enabled
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     : IGMPSnooping SUBMGMT RING
Sync Admin State        : Up
Sync Oper State         : Up
```

```
DB Sync State           : inSync
Num Entries             : 12
Lcl Deleted Entries     : 0
Alarm Entries           : 0
Rem Num Entries         : 12
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
===============================================================================
MCS Application Stats
===============================================================================
Application             : igmp
Num Entries             : 1
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 1
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
```

Notice that the MFIB on PE-2 has also been updated:

```
A:PE-2# show service id 4 mfib
===============================================================================
Multicast FIB, Service 4
===============================================================================
Source Address  Group Address         Sap/Sdp Id               Svc Id   Fwd/Blk
-------------------------------------------------------------------------------
*               *                     sdp:21:4                 Local    Fwd
                                      sdp:24:4                 Local    Fwd
*               225.1.1.1             sap:1/1/2:4.*            Local    Fwd
                                      sdp:21:4                 Local    Fwd
                                      sdp:24:4                 Local    Fwd
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
A:PE-2#
```

The SAP on PE-2 is down to avoid duplicated traffic and loops:

```
A:PE-2# show service id 4 sap 1/1/2:4.*
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 4
SAP                : 1/1/2:4.*               Encap             : qinq
QinQ Dot1p         : Default
Description        : (Not Specified)
Admin State        : Up                      Oper State        : Down
Flags              : StandByForMcRing
Multi Svc Site     : None
Last Status Change : 11/06/2009 18:56:25
Last Mgmt Change   : 11/04/2009 23:43:16
===============================================================================
```

If a failure occurs in the ring-node, the IB-RCC BFD session between PE-1 and PE-2 will go down and the SAP on both PE-1 and PE-2 will be put in operational up state.

Configuration Notes

RNCV (used for ring-node connectivity check) and BFD (used for ring control connection) can either run on the same VLAN in the same IES or VPRN service or can run on different VLAN.

MCS for IGMP or DHCP states on a MC-ring requires ESM since the BSA nodes must know which ring node a subscriber is connected to in case a ring failure occurs. The ring-node name is returned through a Python script, a RADIUS server or through a local user database. This string (int_dest_id) must match one of the ring nodes defined in the redundancy configuration.

Convergence time after a ring failure should be 3 * BFD timer + MCS convergence time. Convergence time after a BSA failure should be likewise.

Note that a debounce timer runs on the MC-ring peers. After a ring failure, the MC-ring converges immediately (after the BFD session times out) and the debounce timer is started. After the ring is fixed and the BFD session is up the MC-ring converges immediately again. If another ring failure occurs before the debounce timer expires, convergence will be slowed down by two (2) seconds. If a third ring failure occurs before the debounce timer expires, four second delays are introduced. In case of a fourth failure, an eight second delay is introduced. 200 seconds of delay is the maximum. The debounce timer can be configured under the MC-ring.

Conclusion

This chapter covers an extension of dual homing support in TPSDA networks based on Layer 2 CO model. The extension addresses networks where multiple access nodes (for example, DSLAMs) are connected in a single ring. The examples show the use of a ring with four access nodes in a ring. The behavior is described in normal operation and in case a failure occurs in the access ring.

Python Cache Support for ESM Applications

This chapter provides information about Python cache support for ESM applications.

Topics in this chapter include:

Applicability

This feature is applicable to 7750 SR-7/12/12e systems, and 7450 systems in mixed mode, with CPM3 or later. It is also applicable to the virtualized simulator but only when running as a distributed simulator. It is not applicable to the 7750 SR c4/12.

The configuration was tested with SR OS release 12.0.R4.

Overview

SR OS sports an embedded Python scripting engine which can be used to manipulate selected messages of protocols including DHCPv4, DHCPv6, RADIUS and Diameter.

Python cache provides a central key-value memory cache with a set of APIs allowing Python scripts to store and retrieve strings across different runs of the same or even different Python scripts.

The following are some basic concepts of Python cache:

- Multiple Python policies can be defined, and each Python policy has a separate cache; only scripts configured in the Python policy can access and share that policy’s cache.
- A cache entry consists of a key and a value, both of them being strings. The key is used to search and fetch the cache entry.

- A cache entry has a lifetime and the system automatically removes expired entries.

The following are the Python cache APIs:

- alc.cache.save(val,key): Save the value identified by the key into the cache.
- alc.cache.retrieve(key): Return the cached entry's value identified by the key.
- alc.cache.clear(key): Remove the entry identified by the key from the cache.
- cache.get_lifetime(key): The system returns an integer indicating the remaining lifetime of the specified entry expressed in seconds.
- cache.set_lifetime(key,new_lifetime): new_lifetime is an integer; the system sets the remaining lifetime of the specified entry (in seconds).

Configuration

The test topology is shown in Figure 205.

Figure 205 **Test Topology**

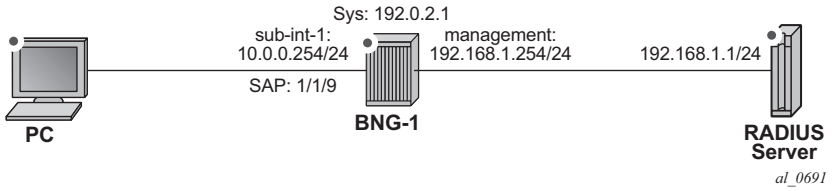

The example shows how the Python cache can be used to store multiple class attributes from the RADIUS access-accept packets and reflect them into RADIUS accounting request packets (start/interim-update/stop).

Test setup:

- A PC is used as DHCPv4 host, connect to SAP 1/1/9 of BNG-1
- SAP 1/1/9 is attached to group interface grp-int-1, which is under subscriber interface sub-int-1 of IES 1
- DHCP host is authenticated via RADIUS server, which resides in management routing instance of BNG-1
- A Python script python-script-1 stores all class attributes in access-accept and add stored attributes into RADIUS accounting requests.
- RADIUS user-name is used as cache key, which is the MAC address of the PC.

The Python cache configuration commands are shown below.

```
config>python>py-policy>
    [no] cache [create]
           [no] entry-size <size>
           [no] max-entries <count>
           [no] max-entry-lifetime [days <days>] [hrs <hours>]
                                   [min <minutes>] [sec <seconds>]
           [no] mcs-peer <ip-address> sync-tag <[32 chars max]>
           [no] minimum-lifetimes
               [no] high-availability <seconds>
               [no] multi-chassis-redundancy <seconds>
               [no] persistence <seconds>
           [no] persistence
           [no] shutdown



config>system>persistence>
    python-policy-cache
        [no] description <desc>
        [no] location <cflash-id>

config>redundancy>multi-chassis>peer>sync>
    [no] python
```

Refer to the SR OS Triple Play Guide for details of above commands. The basic configuration of the Python cache is the **cache create** statement in the python-policy as shown below:

```
config>python>py-policy>
-----------------------------------------------
            cache create
                no shutdown
            exit
            radius access-accept direction ingress script "python-script-1"
            radius accounting-request direction egress script "python-script-1"
-----------------------------------------------
```

The **cache create** configuration enables the cache support for the Python policy. The system's behavior can be tuned in the following aspects:

- Configure **entry-size** and **max-entries** to limit the memory usage.
- Configure **max-entry-lifetime** to specify the maximum lifetime.
- Enable **persistence** to make cache entries persistent across reboot.
- Configure **mcs-peer** to enable Multi-Chassis Synchronization.

In this example only the basic cache configuration is used.

Step 0. Configuring ESM

As a prerequisite ESM must be enabled and as such BNG-1 is configured as follows:

- An **authentication-policy radius-auth-policy-1** is used to authenticate DHCPv4 hosts on group interface grp-int-1.
- A **radius-accounting-policy radius-acct-policy-1** is configured in the **sub-profile sub-profile-1** to enable RADIUS accounting.
- The **user-name** is included in **radius-acct-policy-1** since the User-Name is used as cache entry key.
- Interim-update is enabled in the **radius-acct-policy-1** with an interval 5 minutes.
- The local DHCPv4 server **dhcpv4-svr** is used to assign address to hosts.

```
#--------------------------------------------------
echo "Management Router Configuration"
#--------------------------------------------------
    router management
        radius-server
            server "radius-svr-1" address 192.168.1.1 secret
                                  "iaCuILBunKirJurE4jK2URAnzip6nK32" hash2 create
            exit
        exit
    exit
#--------------------------------------------------
echo "Router (Network Side) Configuration"
#--------------------------------------------------
    router
        dhcp
            local-dhcp-server "dhcpv4-svr" create
            exit
        exit
        interface "system"
            address 192.0.2.1/32
            local-dhcp-server "dhcpv4-svr"
            no shutdown
        exit
    exit

#--------------------------------------------------
echo "Subscriber-mgmt Configuration"
#--------------------------------------------------
    subscriber-mgmt
        authentication-policy "radius-auth-policy-1" create
            password "mcgLj0q0695Dp.pD5DthrCv9Bu8X2qPVSvGYWQmCgUg" hash2
            radius-server-policy "radius-svr-policy-1"
        exit
        radius-accounting-policy "radius-acct-policy-1" create
            update-interval 5
            update-interval-jitter absolute 0
            include-radius-attribute
                user-name
            exit
            radius-server-policy "radius-svr-policy-1"
        exit
```

```
        sla-profile "sla-profile-1" create
        exit
        sub-profile "sub-profile-1" create
            radius-accounting-policy "radius-acct-policy-1"
        exit
    exit
#--------------------------------------------------
echo "Service Configuration"
#--------------------------------------------------
    service
        ies 1 customer 1 create
            subscriber-interface "sub-int-1" create
                address 10.0.0.254/24
                group-interface "grp-int-1" create
                    dhcp
                        server 192.0.2.1
                        gi-address 10.0.0.254
                        no shutdown
                    exit
                    authentication-policy "radius-auth-policy-1"
                    sap 1/1/9 create
                        sub-sla-mgmt
                            def-sub-id use-auto-id
                            def-sub-profile "sub-profile-1"
                            def-sla-profile "sla-profile-1"
                            no shutdown
                        exit
                    exit
                exit
            exit
            no shutdown
        exit
    exit
#--------------------------------------------------
echo "Local DHCP Server (Base Router) Configuration"
#--------------------------------------------------
    router
        dhcp
            local-dhcp-server "dhcpv4-svr" create
                use-gi-address
                pool "addr-pool-1" create
                    subnet 10.0.0.0/24 create
                        options
                            subnet-mask 255.255.255.0
                            default-router 10.0.0.254
                        exit
                        address-range 10.0.0.1 10.0.0.100
                    exit
                exit
                no shutdown
            exit
        exit
    exit
#--------------------------------------------------
echo "AAA Configuration"
#--------------------------------------------------
    aaa
        radius-server-policy "radius-svr-policy-1" create
            servers
```

```
            router "management"
            server 1 name "radius-svr-1"
        exit
    exit
exit
```

Step 1. Create the Python script file.

A Python script is created and stored on the local storage, for example as CF3:\python_cache.py. This script handles the RADIUS packets listed below:

- Access-Accept — All class attributes from the access-accept packets are stored in the cache by combining them into a single string. The user-name (MAC address) is used as the key, and the format of this string is:
  - 1st byte is the number of class attributes in this string
  - 2nd – nth bytes: each byte holds the number of bytes for class-n attributes
  - Rest of bytes: combined string of class-1 to class-n
- Acct-Start — Retrieves the stored combined class string using the user name as key. The combined string is parsed and split into the individual class attributes and then inserted into the packet.
- Interim-Update — The cached entry is parsed and the stored class attributes are inserted, then the lifetime is reset to 15 minutes. This is greater than the interim-update interval (5 minutes) so the cached entry will not expire before next interim-update arrives.
- Acct-Stop — The cached entry is parsed and the stored class attributes are inserted, then the cached entry is removed.

Note: There is some error checking and exception handling logic in the script which causes the script to drop the packet if certain errors occur. This script is only an overview example so the error handling logic should be added according to real application requirements. In addition to the exception handling logic included in the script, the **action-on-fail** command could be used in the **python-script** command to define system's action upon failed execution, for example when an un-captured exception is encountered.

```
#--------------------------------------------------------------------------------
# Name:       Nokia SR OS Python Cache Support Example
# Purpose:    This script is used to demonstrate SR OS python cache support for
#               the following use case:
#                  RADIUS sever returns multiple Class attributes in
#                  access-accept, these Class attributes need to be reflected
#                  in accounting request packets
#--------------------------------------------------------------------------------
from alc import cache
from alc import radius
```

```
import struct
def main():
    radius_header = radius.header()
    if radius_header['code'] == 2: # in case of access-accept
        entry_key = radius.attributes.get(1) # use user-name as the cache entry
                                             # key
        if entry_key == "": #drop the packet if there is no user-name present
            radius.drop()
            return
        class_list = radius.attributes.getTuple(25) # get a list of Class
                                                    # attributes
        if class_list == (): #drop the packet if there is no class present
            radius.drop()
            return
        class_len_str = ''
        class_str = ''
        class_count = 0
        for radius_class in class_list:
            class_len_str += chr(len(radius_class))
            class_str += radius_class
            class_count += 1
        entry_val = chr(class_count)+class_len_str+class_str
        try:
            cache.save(entry_val,entry_key)
        except:
            radius.drop() #drop the packet if cache.save fails
            return

    elif radius_header['code'] == 4: # in case of acct-request
        entry_key = radius.attributes.get(1)
        if entry_key == "": #drop the packet if there is no user-name present
            radius.drop()
            return
        try:
            entry_val = cache.retrieve(entry_key)
        except:#drop the packet if cache.retrieve fails
            radius.drop()
            return
        class_count = ord(entry_val[0])
        pos = class_count+1
        class_list = []
        for i in range(class_count):
            class_len = ord(entry_val[i+1])
            class_list.append(entry_val[pos:pos+class_len])
            pos += class_len
        radius.attributes.set(25,tuple(class_list))
        acct_type=struct.unpack('>I',radius.attributes.get(40))[0]
        if acct_type == 2: # in case of acct-stop
            cache.clear(entry_key) # remove the cache entry
        elif acct_type == 3: # in case of interim-update
            try:
                cache.set_lifetime(entry_key,900) # reset the lifetime
            except:#drop the packet if cache.set_lifetime fails
                radius.drop()
                return

main()
```

Step 2. Configure **python-script** and **python-policy**.

- Enable **python-script-1** to process received (ingress) access-accept packets and transmitted (egress) accounting-request packets.
- Enable the Python cache by configuring **cache create** in **python-policy-1.**
- Reference **python-policy-1** in the **radius-server-policy-1**.

```
#------------------------------------------------
echo "PYTHON Configuration"
#------------------------------------------------
    python
        python-script "python-script-1" create
            primary-url "cf3:/python_cache.py"
            no shutdown
        exit
        python-policy "python-policy-1" create
            cache create
                no shutdown
            exit
            radius access-accept direction ingress script "python-script-1"
            radius accounting-request direction egress script "python-script-1"
        exit
exit
#------------------------------------------------
echo "AAA Configuration"
#------------------------------------------------
    aaa
        radius-server-policy "radius-svr-policy-1" create
            python-policy "python-policy-1"
```

Step 3. Configure the RADIUS server so that:

- The DHCP host is authenticated via its MAC address.
- The RADIUS server returns the following class attributes and values in access-accept packets:
  - “Class-1”
  - “Class-22”
  - “Class-333”

Step 4. Enable debug on BNG-1 to observe the Python cache in action.

- Enable the following debug on BNG-1:

```
debug
    router "Base"
        ip
            dhcp
                detail-level low
                mode egr-ingr-and-dropped
            exit
        exit
    exit
```

```
    router "management"
        radius
            packet-type authentication accounting coa
            detail-level medium
        exit
    exit
    python
        python-script "python-script-1"
            script-all-info
        exit
    exit
exit

A:BNG-1>config>log# info
----------------------------------------------
        log-id 10
            from debug-trace
            to session
        exit
----------------------------------------------
```

Step 5. Initiate DHCPv4 on the PC.

- Initiate the DHCPv4 process on the PC. When the PC sends out a DHCPv4 discover message, BNG-1 contacts the RADIUS server to authenticate the user and a DHCPv4 ESM host is created.
- RADIUS accounting-start will be sent upon host creation and interim-update will be sent every 5 minutes.

The following describes the debug output:

- The system initiates RADIUS authentication upon receipt of a DHCP discovery message.
- The RADIUS server returns an access-accept with the three class attributes.
- The system executes python-script-1, stores the three class attributes.
- After the DHCP host is created, the system sends RADIUS accounting-start and interim-update messages in which python-script-1 adds the three stored class attributes.

```
26 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 3 (grp-int-1),
   received DHCP Boot Request on Interface grp-int-1 (1/1/9) Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:20:fc:1e:cd:53    xid: 0x1e866b74
"

27 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Transmit
  Access-Request(1) 192.168.1.1:1812 id 7 len 63 vrid 4095 pol radius-svr-policy -1
```

```
    USER NAME [1] 17 00:20:fc:1e:cd:53
    PASSWORD [2] 16 R/mc867ChKkdx50PJauB5U
    NAS IP ADDRESS [4] 4 192.168.1.254
"

28 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Receive
  Access-Accept(2) id 7 len 69 from 192.168.1.1:1812 vrid 4095 pol radius-svr-po
licy-1
    CLASS [25] 7 0x436c6173732d31
    CLASS [25] 8 0x436c6173732d3232
    CLASS [25] 9 0x436c6173732d333333
    USER NAME [1] 17 00:20:fc:1e:cd:53
"

29 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 Base Python Output
"Python Output: python-script-1
"

30 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 Base Python Result
"Python Result: python-script-1
"

31 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Script
Access-Accept(2) id 7 len 69 from 192.168.1.1:1812 policy python-policy-
1 stat us success
"

32 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
   transmitted DHCP Boot Request to 192.0.2.1 Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 10.0.0.254
   chaddr: 00:20:fc:1e:cd:53    xid: 0x1e866b74
"

33 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
   received DHCP Boot Reply on 192.0.2.1 Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.0.0.1
   siaddr: 192.0.2.1         giaddr: 10.0.0.254
   chaddr: 00:20:fc:1e:cd:53    xid: 0x1e866b74
"

34 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 3 (grp-int-1),
   transmitted DHCP Boot Reply to Interface grp-int-1 (1/1/9) Port 68

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.0.0.1
   siaddr: 192.0.2.1         giaddr: 10.0.0.254
```

```
   chaddr: 00:20:fc:1e:cd:53    xid: 0x1e866b74
"

35 2014/06/26 03:02:18.34 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 3 (grp-int-1),
   received DHCP Boot Request on Interface grp-int-1 (1/1/9) Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:20:fc:1e:cd:53    xid: 0x1e866b74
"

36 2014/06/26 03:02:18.35 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
   transmitted DHCP Boot Request to 192.0.2.1 Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 10.0.0.254
   chaddr: 00:20:fc:1e:cd:53    xid: 0x1e866b74
"

37 2014/06/26 03:02:18.35 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base),
   received DHCP Boot Reply on 192.0.2.1 Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.0.0.1
   siaddr: 192.0.2.1         giaddr: 10.0.0.254
   chaddr: 00:20:fc:1e:cd:53    xid: 0x1e866b74
"

38 2014/06/26 03:02:18.38 UTC MINOR: DEBUG #2001 Base PIP
"PIP: DHCP
instance 1 (Base), interface index 3 (grp-int-1),
   transmitted DHCP Boot Reply to Interface grp-int-1 (1/1/9) Port 68

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.0.0.1
   siaddr: 192.0.2.1         giaddr: 10.0.0.254
   chaddr: 00:20:fc:1e:cd:53    xid: 0x1e866b74
"

39 2014/06/26 03:02:18.38 UTC MINOR: DEBUG #2001 Base Python Output
"Python Output: python-script-1
"

40 2014/06/26 03:02:18.38 UTC MINOR: DEBUG #2001 Base Python Result
"Python Result: python-script-1
RADIUS Attribute: Type 25, SET
    'Class-1'
    'Class-22'
    'Class-333'
"
```

```
41 2014/06/26 03:02:18.38 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Transmit
  Accounting-Request(4) 192.168.1.1:1813 id 8 len 152 vrid 4095 pol radius-svr-
policy-1
    STATUS TYPE [40] 4 Start(1)
    NAS IP ADDRESS [4] 4 192.168.1.254
    USER NAME [1] 17 00:20:fc:1e:cd:53
    SESSION ID [44] 63 00:20:fc:1e:cd:53|1/1/9@1/1/9@sla-profile-1_2014/06/26 03
:02:18
    EVENT TIMESTAMP [55] 4 1403751738
    CLASS [25] 7 0x436c6173732d31
    CLASS [25] 8 0x436c6173732d3232
    CLASS [25] 9 0x436c6173732d333333
"

42 2014/06/26 03:02:18.38 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Receive
  Accounting-Response(5) id 8 len 20 from 192.168.1.1:1813 vrid 4095 pol radius-svr-
policy-1
"

43 2014/06/26 03:02:18.38 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Script
  Accounting-Response(5) id 8 len 20 from 192.168.1.1:1813 policy python-policy-
1 status success
"
44 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 Base Python Output
"Python Output: python-script-1
"

45 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 Base Python Result
"Python Result: python-script-1
RADIUS Attribute: Type 25, SET
    'Class-1'
    'Class-22'
    'Class-333'
"

46 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Transmit
  Accounting-Request(4) 192.168.1.1:1813 id 9 len 302 vrid 4095 pol radius-svr-
policy-1
    STATUS TYPE [40] 4 Interim-Update(3)
    NAS IP ADDRESS [4] 4 192.168.1.254
    USER NAME [1] 17 00:20:fc:1e:cd:53
    SESSION ID [44] 63 00:20:fc:1e:cd:53|1/1/9@1/1/9@sla-profile-1_2014/06/26 03
:02:18
    SESSION TIME [46] 4 300
    EVENT TIMESTAMP [55] 4 1403752038
    VSA [26] 12 Alcatel(6527)
      INPUT_INPROF_OCTETS_64 [19] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      INPUT_OUTPROF_OCTETS_64 [20] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      INPUT_INPROF_PACKETS_64 [23] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      INPUT_OUTPROF_PACKETS_64 [24] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      OUTPUT_INPROF_OCTETS_64 [21] 10 0x00010000000000000bea
```

```
    VSA [26] 12 Alcatel(6527)
      OUTPUT_OUTPROF_OCTETS_64 [22] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      OUTPUT_INPROF_PACKETS_64 [25] 10 0x00010000000000000019
    VSA [26] 12 Alcatel(6527)
      OUTPUT_OUTPROF_PACKETS_64 [26] 10 0x00010000000000000000
    CLASS [25] 7 0x436c6173732d31
    CLASS [25] 8 0x436c6173732d3232
    CLASS [25] 9 0x436c6173732d333333
"

47 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Receive
  Accounting-Response(5) id 9 len 20 from 192.168.1.1:1813 vrid 4095 pol radius-svr-
policy-1
"


48 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Script
  Accounting-Response(5) id 9 len 20 from 192.168.1.1:1813 policy python-policy-
1 status success
"
44 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 Base Python Output
"Python Output: python-script-1
"

45 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 Base Python Result
"Python Result: python-script-1
RADIUS Attribute: Type 25, SET
    'Class-1'
    'Class-22'
    'Class-333'
"

46 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Transmit
  Accounting-Request(4) 192.168.1.1:1813 id 9 len 302 vrid 4095 pol radius-svr-
policy-1
    STATUS TYPE [40] 4 Interim-Update(3)
    NAS IP ADDRESS [4] 4 192.168.1.254
    USER NAME [1] 17 00:20:fc:1e:cd:53
    SESSION ID [44] 63 00:20:fc:1e:cd:53|1/1/9@1/1/9@sla-profile-1_2014/06/26 03
:02:18
    SESSION TIME [46] 4 300
    EVENT TIMESTAMP [55] 4 1403752038
    VSA [26] 12 Alcatel(6527)
      INPUT_INPROF_OCTETS_64 [19] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      INPUT_OUTPROF_OCTETS_64 [20] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      INPUT_INPROF_PACKETS_64 [23] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      INPUT_OUTPROF_PACKETS_64 [24] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      OUTPUT_INPROF_OCTETS_64 [21] 10 0x0001000000000000bea
    VSA [26] 12 Alcatel(6527)
      OUTPUT_OUTPROF_OCTETS_64 [22] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
```

```
      OUTPUT_INPROF_PACKETS_64 [25] 10 0x00010000000000000019
    VSA [26] 12 Alcatel(6527)
      OUTPUT_OUTPROF_PACKETS_64 [26] 10 0x00010000000000000000
    CLASS [25] 7 0x436c6173732d31
    CLASS [25] 8 0x436c6173732d3232
    CLASS [25] 9 0x436c6173732d333333
"

47 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Receive
  Accounting-Response(5) id 9 len 20 from 192.168.1.1:1813 vrid 4095 pol radius-svr-
policy-1
"

48 2014/06/26 03:07:18.71 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Script
  Accounting-Response(5) id 9 len 20 from 192.168.1.1:1813 policy python-policy-
1 status success
"
```

As the debug output shows, **python-script-1** stores the three class attributes from the access-accept message which then are reflected into the accounting-start and interim-update messages.

The **tools dump python python-policy** <*name*> **cache** command can be used to show the existing cached entries in the specified python-policy:

```
A:BNG-1# tools dump python python-policy "python-policy-1" cache
===============================================================================
Python policy cache "python-policy-1" entries
===============================================================================
Key       : 00:20:fc:1e:cd:53
Value     : (hex) 03 07 08 09 43 6c 61 73 73 2d 31 43 6c 61 73 73 2d 32 32 43 6c 61
73 73 2d 33 33 33
Time Left : 0d 00:12:08
DDP Key   : N/A
===============================================================================
```

Step 6. Release DHCPv4 lease on the PC

- Release DHCPv4 lease on the PC which is sent to BNG-1.
- DHCPv4 release message from the PC will trigger BNG-1 to remove the ESM host on BNG-1.
- BNG-1 will send an accounting-stop packet to the RADIUS server.

The following is the debug output:

```
"Python Output: python-script-1
"

67 2014/06/26 03:26:14.84 UTC MINOR: DEBUG #2001 Base Python Result
"Python Result: python-script-1
RADIUS Attribute: Type 25, SET
    'Class-1'
    'Class-22'
    'Class-333'
"
```

```
68 2014/06/26 03:26:14.85 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Transmit
  Accounting-Request(4) 192.168.1.1:1813 id 13 len 308 vrid 4095 pol radius-svr-
policy-1
    STATUS TYPE [40] 4 Stop(2)
    NAS IP ADDRESS [4] 4 192.168.1.254
    USER NAME [1] 17 00:20:fc:1e:cd:53
    SESSION ID [44] 63 00:20:fc:1e:cd:53|1/1/9@1/1/9@sla-profile-1_2014/06/26 03
:02:18
    SESSION TIME [46] 4 1436
    TERMINATE CAUSE [49] 4 User Request(1)
    EVENT TIMESTAMP [55] 4 1403753174
    VSA [26] 12 Alcatel(6527)
      INPUT_INPROF_OCTETS_64 [19] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      INPUT_OUTPROF_OCTETS_64 [20] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      INPUT_INPROF_PACKETS_64 [23] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      INPUT_OUTPROF_PACKETS_64 [24] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      OUTPUT_INPROF_OCTETS_64 [21] 10 0x0001000000000001a36
    VSA [26] 12 Alcatel(6527)
      OUTPUT_OUTPROF_OCTETS_64 [22] 10 0x00010000000000000000
    VSA [26] 12 Alcatel(6527)
      OUTPUT_INPROF_PACKETS_64 [25] 10 0x00010000000000000037
    VSA [26] 12 Alcatel(6527)
      OUTPUT_OUTPROF_PACKETS_64 [26] 10 0x00010000000000000000
    CLASS [25] 7 0x436c6173732d31
    CLASS [25] 8 0x436c6173732d3232
    CLASS [25] 9 0x436c6173732d333333
"

69 2014/06/26 03:26:14.85 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Receive
  Accounting-Response(5) id 13 len 20 from 192.168.1.1:1813 vrid 4095 pol radius
-svr-policy-1
"

70 2014/06/26 03:26:14.85 UTC MINOR: DEBUG #2001 management RADIUS
"RADIUS: Script
  Accounting-Response(5) id 13 len 20 from 192.168.1.1:1813 policy python-policy
-1 status success
"
```

As the debug output shows, **python-script-1** inserts three RADIUS class attributes in the accounting-stop message.

The **tools dump python python-policy** <*name*> **cache** command can be used to verify the cached entry has been removed:

```
A:BNG-1# tools dump python python-policy "python-policy-1" cache
===============================================================================
Python policy cache "python-policy-1" entries
===============================================================================
===============================================================================
```

Step 7. Manually change the lifetime of a cached entry (optional).

A **tools** command can be used to manually change the lifetime of an existing cached entry, with following syntax:

```
tool perform python-policy <name> cache {hex-key <hex-str>|string-key <str>} set-
lifetime <newlifetime>
```

However, manually changing the lifetime might cause issues with the Python script (for example, if reducing the lifetime causes the entry to expire) that needs the cached entry so it should be used with caution.

To demonstrate this recreate the cached entry by initiating a DHCPv4 discover from the PC (see Step 5, Initiate DHCPv4 on the PC.), then change the lifetime to 20 minutes.

```
A:BNG-1# tools dump python python-policy "python-policy-1" cache
===============================================================================
Python policy cache "python-policy-1" entries
===============================================================================
Key       : 00:20:fc:1e:cd:53
Value     : (hex) 03 07 08 09 43 6c 61 73 73 2d 31 43 6c 61 73 73 2d 32 32 43 6c 61
73 73 2d 33 33 33
Time Left : 0d 00:09:48
DDP Key   : N/A
===============================================================================
A:BNG-1# tools perform python-policy "python-policy-1" cache string-
key "00:20:fc:1e:cd:53" set-lifetime 1200
A:BNG-1# tools dump python python-policy "python-policy-
1" cache
===============================================================================
Python policy cache "python-policy-1" entries
===============================================================================
Key       : 00:20:fc:1e:cd:53
Value     : (hex) 03 07 08 09 43 6c 61 73 73 2d 31 43 6c 61 73 73 2d 32 32 43 6c 61
73 73 2d 33 33 33
Time Left : 0d 00:19:57
DDP Key   : N/A
===============================================================================
```

Step 8. Manually remove a cached entry (optional).

Manually removing an existing cached entry can be done using a clear command with following syntax:

```
clear python python-policy <name> cache {hex-key <hex-str>|string-key <str>}
```

Manually removing a cached entry can result in unexpected results (for example, if a script expects an entry to exist but it has been removed), so this should be used with caution.

The following command sequence demonstrates the effect of the clear command after initiating a DHCPv4 discover from the PC to recreate the cached entry (see Step 5, Initiate DHCPv4 on the PC.).

```
A:BNG-1# tools dump python python-policy "python-policy-
```

```
1" cache
===============================================================================
Python policy cache "python-policy-1" entries
===============================================================================
Key       : 00:20:fc:1e:cd:53
Value     : (hex) 03 07 08 09 43 6c 61 73 73 2d 31 43 6c 61 73 73 2d 32 32 43 6c 61
73 73 2d 33 33 33
Time Left : 0d 00:11:39
DDP Key   : N/A
===============================================================================
A:BNG-1# clear python python-policy "python-policy-1" cache string-
key "00:20:fc:1e:cd:53"
A:BNG-1# tools dump python python-policy "python-policy-
1" cache

===============================================================================
Python policy cache "python-policy-1" entries
===============================================================================
===============================================================================
```

Configuration and Operational Guidelines

The following is a list of configuration and operational guidelines that a user should follow when using the Python cache:

- SR OS has a limit on the total amount of memory used for the python cache since the python cache can be demanding with respect to memory usage. The maximum memory allocated for the cache system wide is restricted to 256MB. However, it is good practice to configure per python-policy limits using the **entry-size** and **max-entries** commands; by doing this, one python-policy's cache will not impact another python-policy's cache memory usage.
- For applications needing a cache entry for the entire lifetime of an ESM host, lifetime management is essential. If the lifetime is too long then unneeded entries might reside in the system, wasting memory; if the lifetime is too short then entries might expire while they are still needed. One way to address this is by initially setting a relative short lifetime and then using the RADIUS interim-update message as a trigger to reset a new lifetime. This new lifetime should be larger than the interim-update interval. Then the entry should be removed by a script when an accounting-stop message is sent.
- With MCS enabled, each python cached entry will have a corresponding MCS record, resulting in each python cache entry consuming twice amount of memory. For example, 256MB cached entries would consume additional 256MB memory for MCS records.

- Choosing the right key is important, a network designer needs to choose the key meeting the application requirement in terms of their uniqueness. For example: if an application needs to store per-host information then the key for that host must be unique (for example, its MAC address or remote-id). The key must be derived from the trigger packet. For example: if an application needs to store information on DHCP discovery and retrieve it on receiving a RADIUS accounting-request message for the same host, then the script needs to be able derive the same key from both the DHCP discovery as well as from the RADIUS accounting-request message.
- Using tools or clear commands to manually change cached entries could cause problems if the entries are needed by a script. Only use these commands when it is absolutely necessary.
- The minimum-lifetimes exist to make the system more efficient in handling system memory before a cache entry could be synced or made persistent (e.g. when a new entry is created or when MCS is enabled). The cache entry's lifetime must be equal or larger than the configured minimal-lifetime listed below for that function to occur.
  1. high-availability — The minimum lifetime of a cache entry for it to be synchronized from the active CPM to the standby CPM.

     The default is 0 seconds, resulting in all entries being synchronized.
  2. multi-chassis-redundancy — The minimum lifetime of a cache entry for it to be synchronized between chassis.

     The default is 0 seconds, resulting in all entries being synchronized.
  3. persistence — The minimum lifetime of a cache entry for it to be written to the persistency file.

     The default is 0 seconds, resulting in all entries will be synchronized.

Conclusion

The Python cache provides a very powerful and flexible way to share information across different Python scripts in SR OS.

RADIUS-Triggered Dynamic Data Service Provisioning

This chapter describes advanced RADIUS-triggered dynamic data service provisioning configurations.

Topics in this chapter include:

Applicability

This example is applicable to all 7750 SR and 7450 ESS in mixed mode with multi-core CPM (CPM-3 and later) and ESM capability.

This feature is not supported on the 7950 XRS or 7450 ESS.

The configuration was tested on release 11.0.R2.

Overview

RADIUS-triggered dynamic data services enables a zero touch, single-ended provisioning model for business services on the basis of Enhanced Subscriber Management functionality.

Triggered by the authentication of a single or dual stack PPPoE or IPoE session or a single stack IPv4 host as the “control channel” from the business CPE, parameters are passed in a RADIUS Access Accept or Change of Authorization (CoA) message to set up one or multiple Layer 2 or Layer 3 data services.

This concept removes the need to have an Operations Support System (OSS) responsible for the service provisioning and is particularly beneficial in a highly dynamic network environment, where physical network topologies – especially in the access – change frequently. With a regular service provisioning, frequent changes would be hard to keep track of. In the RADIUS-based model the service gets instantiated wherever it “pops-up” in the network. Even planned customer moves to a different office would not require advanced notifications and lead times but could be instantaneous, assuming the pure physical connectivity is given.

A variation of the current service offering will only require one or a few modified service parameters in the RADIUS user database and does not require timely and costly IT changes (for RADIUS those service parameters are just attributes; the RADIUS server does not check the logic). This speeds up the time-to-market for new service offerings, which is another big advantage.

Taking this logic to its full extent, it becomes immediately clear that the managed business CPE terminating the carrier service (and being responsible for the PPP or DHCP control channel) also needs to be provisioned in the most flexible way. Through the control channel or a dedicated management channel instantiated as the first dynamic data service, the business CPE should get its full configuration from a configuration server via a pre-populated configuration file. The details of the CPE provisioning are outside of the scope of this example and therefore not discussed further.

As the whole approach is centered around the principle of “highly flexible in a highly dynamic environment”, it is naturally required to maintain as little state information about connections in the RADIUS parameter attributes as possible. For example, fixed remote peer IP-addresses for the SDPs used in a VPLS service in the RADIUS parameter lists would remove all the flexibility and would not allow access services to be moved dynamically. As such the data services for this functionality focuses on those types where a control protocol like BGP is used to exchange VPN membership information. Dynamic data services supported include local Epipe VLL services, Epipe VLL services with dynamic Multi-Segment PseudoWires (MS-PWs) (FEC129) or spoke SDP, VPLS services with BGP-AD PWs or mesh/spoke SDPs, IES, and VPRN services.

To display the complete white list of dynamic data services CLI configuration commands, use the **tools dump service dynamic-services command-list** CLI command.

A Python script interface adds a flexible abstraction layer so that only the business user specific service parameters (service type, IP address, QoS and filter parameters, etc.) are required from RADIUS and are then used in a CLI template to set up the target service.

The setup sequence is shown in Figure 206 with the example of a VPLS service.

Figure 206 **Principle Model of Dynamic Data Services**

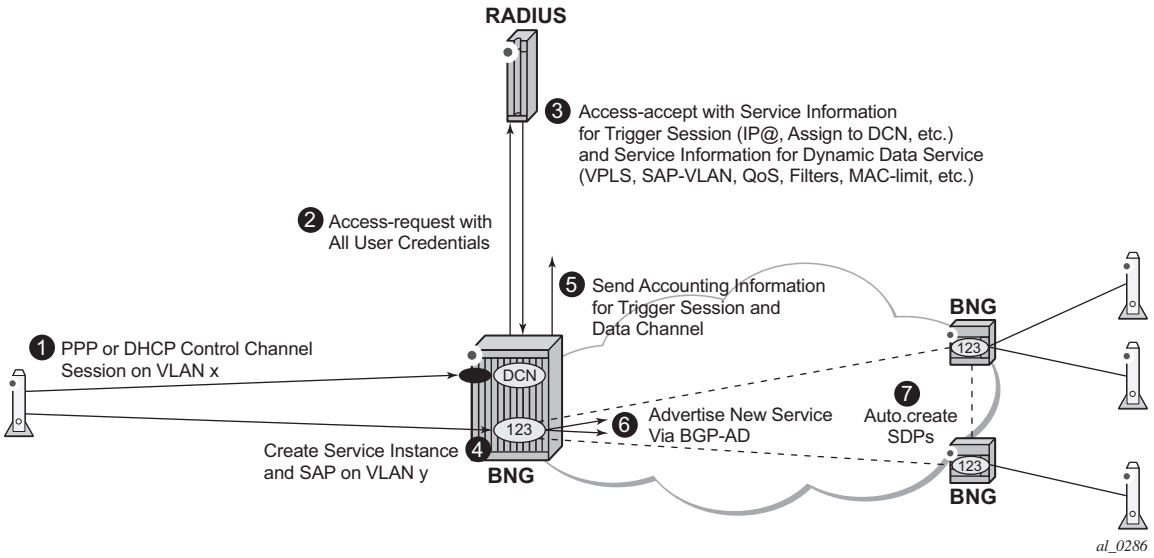

1. Business CPE initiates a PPP or DHCP “control channel” session. This session is important to let the BNG and RADIUS understand the existence of a new access circuit and the location of the service endpoint.
2. BNG sends an Access-Request with all user credentials to RADIUS.
3. RADIUS replies using an Access-Accept with attributes for the PPP/DHCP control channel and attributes for the dynamic data service (Service-type, SAP-VLAN, QoS, Filter, etc.).
4. BNG creates a dynamic data service instance (if it is first access-circuit for this service) and a SAP, thus completing the service configuration.
5. BNG sends an Accounting Start for PPP/DHCP control channel session and also for the dynamic data service session (and subsequently interim accountings and accounting stop for both).
6. BNG advertises VPLS instance-ID (123) via BGP-AD to other PEs/BNGs.
7. PEs/BNGs with same service instance will auto-establish SDPs to the BNG.

The result is a fully functional service which is the same as a traditionally configured service.

The lifetime of the dynamic data services are bound to the existence of the control channel session. If, for whatever reason, the control channel session is torn down all associated dynamic data services will also be terminated.

Dynamic data service SAPs have to be located on dot1q or qinq encapsulated Ethernet ports and can be part of a LAG.

Both XML accounting and RADIUS accounting can be enabled on a dynamic data service SAP. The RADIUS accounting data can be sent to up to two different RADIUS servers.

There is a strict separation of services created by dynamic service provisioning and services created via the CLI or through other standard mechanisms (5620 SAM, SNMP). It is therefore not allowed to:

- create a dynamic services object in a local provisioned CLI/SNMP context (e.g. create a dynamic SAP in a local provisioned VPLS).
- create a local provisioned object in a dynamic service context (e.g. create a SAP via CLI/SNMP in a dynamic VPLS service).
- change parameters in a local provisioned CLI/SNMP context using the dynamic services model (change system name with dynamic services provisioning).
- change parameters with the CLI/SNMP in a dynamically created context.
- delete a local provisioned object using dynamic provisioning model.
- delete a dynamic provisioning object using the CLI/SNMP.
- create a reference to a dynamic services object in a local provisioned CLI/SNMP context (reference to dynamic interface in **router ospf**)

A special command exists to overcome some of the above rules. This command is designed to ease Python script creation and testing and not for normal operations. This is discussed in Configuration.

Configuration

It is assumed that the reader is familiar with the regular Enhanced Subscriber Management (ESM) functionality as well as with general service related configurations. Furthermore certain knowledge about Python programming is also assumed.

The test topology is shown in Figure 207.

Figure 207 **Test Topology**

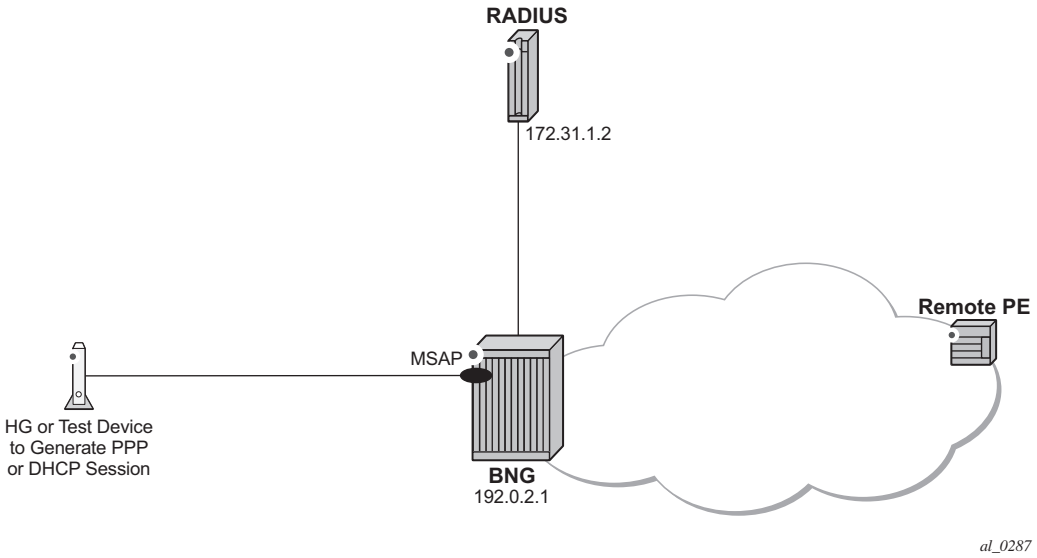

The pure service setup can be tested with a single node acting as BNG. However an Epipe service between two nodes will not normally become status “up” with only one endpoint in an up state. As such, for packets should be sent through the established dynamic data service, a remote PE could also be configured. The remote PE could have its data service configured in a regular fashion, meaning via CLI/SNMP or 5620 SAM.

The required functionality on the BNG is divided into multiple building blocks. The following sections discuss each building block in detail.

Figure 208 **Building Blocks of Dynamic Data Services**

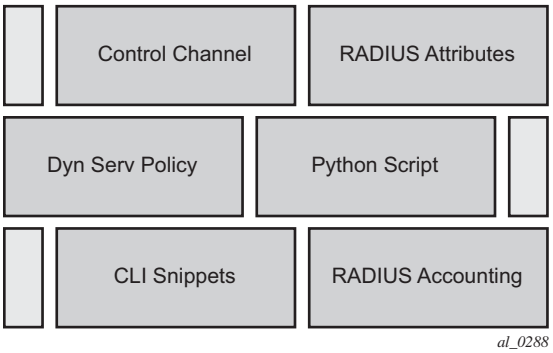

Based on a PPP or DHCP control session, RADIUS will return the required parameters for the dynamic data service via dedicated Vendor Specific Attributes (VSAs). The existence of those attributes in the RADIUS Accept message will trigger the relaying of the parameters relating to those attributes towards the Python script defined in the dynamic service policy, which will process them to generate the regular CLI output for the various service types (IES, VPRN, Epipe, VPLS).

For efficiency and flexibility the Python script needs to be structured into different parts per service which then reference each other internally. Those parts are called snippets.

Finally, as the services are initiated from RADIUS, RADIUS accounting messages per dynamic data service will be sent to the RADIUS server as a necessary feedback mechanism to inform the RADIUS server about a successful or failed service setup.

Building Block: Control Channel

The configuration to authenticate and instantiate a dynamic data service control channel is identical to a residential Enhanced Subscriber Management (ESM) configuration. Examples for this can be found in other chapters of the advanced configuration guide and will not be covered here in detail.

Building Block: Dynamic Services Policy

The dynamic services parameters are configured under the **configure service dynamic-services** CLI context. The following output shows two policy examples.

```
configure service dynamic-services
              dynamic-services-policy "dynamic-services-1" create
                accounting-1
                    server-policy "radius-server-policy-1"
                    update-interval min 5
                exit
                accounting-2
                    server-policy "radius-server-policy-2"
                    stats-type time
                    update-interval min 5
                    update-interval-jitter absolute 10
                exit
                cli-user "dynuser"
                description "Dynamic Service Policy #1"
                sap-limit 4000
                script-policy "script-policy-1"
            exit
            dynamic-services-policy "dynamic-services-2" create
                accounting-1
```

```
            server-policy "radius-server-policy-2"
        stats-type volume-time
            update-interval min 30
            update-interval-jitter absolute 20
        exit
        accounting-2
            server-policy "radius-server-policy-2"
            stats-type time
            update-interval min 5
            update-interval-jitter absolute 10
        exit
        cli-user "dynuser"
        description "Dynamic Service Policy #2"
        sap-limit 100
        script-policy "script-policy-2"
    exit
    service-range 1000 10000
    timers
        setup-timeout access-accept 3
    exit
```

Details of each command and the possible parameters can be found in the SR OS Triple Play Guide in the RADIUS Triggered Dynamic Data Services section.

On the top command level under the dynamic-services sub-tree there are three options:

- dynamic-services-policy
- service-range
- timers

The setup-timeout value under **timers** is used to limit the maximum delay allowed for a dynamic data service setup. In addition, it also protects the node during times where there is a high load on the CPU. If a requested dynamic data service cannot be established in the specified time the request will be dropped.

Dynamic data services are not preferred over regular ESM subscribers. As such, given a BNG with a mix of residential ESM subscribers and business customers with dynamic data services, all compete for the same CPU resources to establish the connections.

However, dynamic data services are expected to have a very long lifetime compared to potentially very dynamic lifetimes for residential subscribers. In a regular operating mode the amount of additional setup requests for dynamic data services should be relatively small. Only in the event of a node reboot will all users again compete to gain access, where longer setup-times are inevitable.

The service-range value reserves a certain amount of service IDs for the use of dynamic data services. The configured range is no longer available for regular provisioned services configured via the CLI/SNMP.

The dynamic-services-policy contains a CLI-user identifier, SAP-limits, accounting parameters and reference to a Python script policy which is used when creating a dynamic data service. Multiple dynamic services policies can be created to enable different profiles to be used for different users/customers or services (as an example, two different departments within the service provider, one responsible for Layer 2 services, one for Layer 3 services). The policy used for a dynamic data service is determined from the Alc-Dyn-Serv-Policy [26-6527-167] RADIUS attribute. If the attribute is not present and a policy named **default** exists, then the **default** policy is used, otherwise the dynamic data service creation fails.

Up to two accounting server policies can be defined. This allows the use of separate RADIUS accounting servers independent from the accounting servers used for residential services. The parameters defined in the accounting sections are the default values which are used if no specific values are sent via RADIUS VSAs.

As the service is established via RADIUS, a feedback mechanism towards RADIUS is most likely required which would be at least RADIUS start and stop messages per service/session. In addition performance counters (with a fixed set of parameters) can also be included in the RADIUS messages. It is also possible to use the standard service-accounting under the service instance and remove any counters from the RADIUS accounting messages.

The specification of a CLI user allows linking of the dynamic data service to a specific user-profile. In addition, this facilitates limiting of the scope of allowed service configurations even further, based on the specified context under the user profile.

The CLI-user needs to be configured locally on the node and needs to have a local user profile (remote authorization via TACACS/RADIUS is not possible).

The radius-script-policy is configured under the **configure aaa** CLI context.

```
configure aaa
radius-script-policy "script-policy-2"
    action-on-fail passthrough
            primary
                script-url "cf3:/scripts/dyn_services.py"
                no shutdown
            exit
            secondary
                script-url ftp://user*:*pwd@10.255.137.80/scripts/dyn_services.py"
                no shutdown
            exit
    exit
exit
```

The parameters are no different to what have been defined generally for the use of Python scripting on the BNG.

When the very first session request arrives, the Python script is loaded into memory and executed. For all subsequent session requests the script is executed without the need for a reload. It is possible for both primary and secondary locations to be FTP sites (the small transfer delay for the first session is acceptable), however, it is recommended to have a compact-flash (cf1 or cf2) as the primary location and a remote location as backup.

Building Block: RADIUS Attributes

A series of vendor specific attributes (VSAs) have been defined to setup, teardown or modify dynamic data services from RADIUS.

The VSAs and their meaning are as follows:

- Alc-Dyn-Serv-SAP-Id [26-6527-164], type “string”

  This attribute identifies the dynamic service SAP. The format can be any valid Ethernet SAP format (dot1q or qinq encapsulation), including LAGs. A wildcard (“#”) can be specified for the port field and optionally for one of the tag fields of a qinq interface. To define the dynamic data service SAP-ID, the wildcard fields are replaced with the corresponding field from the Control Channel SAP-ID.

  Examples: “1/2/7:10.100” or “#:#.100”

- Alc-Dyn-Serv-Script-Action [26-6527-166], type “integer”

  A mandatory VSA in a COA to the control channel accounting session ID or the accounting session ID of the dynamic data service (only applicable for modify or teardown). Tells the system what script action is required: setup, modify or teardown of a dynamic data service.

  Values: 1=setup, 2=modify, 3=teardown

- Alc-Dyn-Serv-Policy [26-6527-167], type “string”

  Specifies the dynamic service policy to use for provisioning the dynamic service. The policy must be configured in the “configure service dynamic-services dynamic-services-policy < dynsrv-policy-name>” CLI context.

- Alc-Dyn-Serv-Script-Params [26-6527-165], type “string”

  This VSA contains parameters that can be used by the Python script to setup or modify a dynamic data service. The parameters can be split into multiple instances of the same attribute, linked together by the same tag, that is, the parameters can cross an attribute boundary. The concatenation of all “Alc-Dyn-Serv-Script-Params” attributes with the same tag in a single message must be formatted as “function-key = {dictionary}” where function-key specifies which Python functions will be called and {dictionary} contains the actual parameters in a Python dictionary structure format.

  Example: “business_1 = { 'as_id' : '100', 'comm_id' : '200', 'if_name' : 'itf1', 'ipv4_address' : '172.16.1.1', 'egr_ip_filter' : '100' , 'routes' : [{'to' : '172.16.100.0/24', 'next-hop' : '172.16.1.2'}, {'to' : '172.16.200.0/24', 'next-hop' : '172.16.1.2'}]}

  The above example shows each parameter with a keyword and the associated value. Alternatively only the parameter values can be sent with a pre-defined (and always constant) sequence.

  Example: “business_1 = {“t”: '100', '200', 'itf1', '172.16.1.1', '100', '172.16.100.0/24', '172.16.1.2', '172.16.200.0/24', '172.16.1.2'}.

- Alc-Dyn-Serv-Acct-Interim-Ivl-1 [26-6527-168], type “integer”

  This VSA defines the number of seconds between each accounting interim update message for the primary accounting server. It overrides the local configured “update-interval” value in the dynamic services policy “accounting-1” CLI context. A value of 0 (zero) corresponds to no accounting interim update messages. A value [1..299] seconds is rounded to 300s (min. CLI value) and a value above 15552000 seconds (180 days, maximum CLI value) is rounded to the maximum CLI value.

  Range = 0 | [300 - 15552000].

- Alc-Dyn-Serv-Acct-Interim-Ivl-2 [26-6527-169], type “integer”

  Same function and values as Alc-Dyn-Serv-Acct-Interim-Ivl-1 [26-6527-168], for the second accounting server. It overrides the locally configured “update-interval” value in the dynamic services policy “accounting-2” CLI context.

- Alc-Dyn-Serv-Acct-Stats-Type-1 [26-6527-170], type “integer”

  Enable or disable dynamic data service accounting to the primary accounting server and specify the type of statistics that should be reported: volume and time or time only. It overrides the locally configured value in the dynamic services policy “accounting-1” CLI context.

  Values: 1=off, 2=volume-time, 3=time

- Alc-Dyn-Serv-Acct-Stats-Type-2 [26-6527-171], type “integer”

  Enable or disable dynamic data service accounting to the secondary accounting server and specify the type of statistics that should be reported: volume and time or time only. It overrides the locally configured “stats-type” value in the dynamic services policy “accounting-2” CLI context.

  Values: 1=off, 2=volume-time, 3=time

All VSAs are tagged to enable manipulation of up to 32 (tag values 0..31) dynamic data services in a single RADIUS message. VSAs with an identical tag belong to the same dynamic data service.

The use of the VSAs in RADIUS Access-Accept, CoA and Disconnect Messages is detailed in Table 41. An Access-Accept message can only contain dynamic data service setup requests. A CoA can be used to setup, modify or terminate a dynamic data service. A Disconnect Message can only be used to terminate a dynamic data service.

Table 41 **Dynamic Service Attribute List for Setup, Modify and Teardown**

Attribute Name	Access Accept	CoA			Disc. Message	Comment
	Setup	Setup	Modify	Teardown	Teardown	
Acct-Session-Id	N/A	M	M	M	M	Acct-Session-Id of the Control Channel or in case of a CoA: any other valid CoA key for ESM hosts/sessions.
Alc-Dyn-Serv-SAP-Id	M	M(*)	M(*)	M(*)	N/A	Identifies the dynamic data service
Alc-Dyn-Serv-Script-Params	O	M(*)	M(*)	N/A	N/A	For a Modify, the script parameters represent the new parameters required for the change.
Alc-Dyn-Serv-Script-Action	O	M(*)	M(*)	M(*)	N/A	Must be “setup” if specified in an access-accept.
Alc-Dyn-Serv-Policy	O	O	O	O	N/A	The **default** policy used when not specified for create. In CoA, must be same as used for Setup if Specified for Modify or Teardown.

Table 41 **Dynamic Service Attribute List for Setup, Modify and Teardown (Continued)**

Attribute Name	Access Accept	CoA			Disc. Message	Comment
	Setup	Setup	Modify	Teardown	Teardown	
Alc-Dyn-Serv-Acct-Interim-Ivl-1	O	O	X(**)	X(**)	N/A	
Alc-Dyn-Serv-Acct-Interim-Iv2	O	O	X(**))	X(**))	N/A	
Alc-Dyn-Serv-Acct-Stats-Type-1	O	O	X(**)	X(**))	N/A	
Alc-Dyn-Serv-Acct-Stats-Type-2	O	O	X(**))	X(**))	N/A	
M = Mandatory, O= Optional, X = May not, N/A = Not Applicable (ignored)						
(*) = CoA Nak'd, if not specified (Error Cause: 402 - Missing Attribute)						
(**) = CoA Nak'd if specified (Error Cause:405 - Unsupported Service)						

To summarize, Table 42 shows resulting dynamic service script actions as function of the RADIUS message (Access-Accept, CoA or DM) and the target (Control Channel or Dynamic Service SAP).

Table 42 **Dynamic Service Actions on Control- and Data-Channel**

Target	RADIUS Message	Dynamic Service Script Action	Comments
Control Channel	Access-Accept	Setup	Up to 32 dynamic data services in single message. Alc-Dyn-Serv-Script-Action VSA optional.
		Modify/Teardown	Not supported.
	CoA (acct-session-id or any other valid CoA key for ESM hosts/sessions)	Create/Modify/ Teardown	Cannot be mixed with session/post parameter changes in the same RADIUS message (results in CoA NAK). Up to 32 dynamic data services in single message. Alc-Dyn-Serv-Script-Action VSA mandatory.
	Disconnect	N/A	Teardown the Control Channel session and all associated dynamic data services.

Table 42 **Dynamic Service Actions on Control- and Data-Channel (Continued)**

Target	RADIUS Message	Dynamic Service Script Action	Comments
Dynamic Service	CoA (acct-session-id of the dynamic data service sap)	Modify/Teardown	Only single dynamic data service per message (Acct-Session-Id). Alc-Dyn-Serv-Script-Action VSA mandatory.
		Setup	Not supported.
	Disconnect (acct-session-id of the dynamic data service sap)	N/A	Teardown the corresponding dynamic data service.

Building Block: Python Script

Dynamic data services scripts are built using a Python script engine. The following dedicated functions are available in the alc.dyn module:

- dyn.reference(function-key, reference-id string, dictionary)

  This function creates a dynamic reference to another function in the script. This function eases the creation of N:1 relationships in the script. For more information about use cases, see Building Block: CLI Snippets. The function-key specifies the key in the action dictionary to find the corresponding setup/modify/teardown function calls.

  The reference-id (typically derived from a parameter specified from RADIUS, for example: service-name) specifies a unique instance string that identifies this reference.

  The dictionary specifies a dictionary with parameters that can be used in the parent function to generate CLI script output.

- dyn.action(d)

  When called, the dyn.action will take the “function-key” string specified in the Alc-Dyn-Serv-Script-Params attribute, and perform a lookup in the specified dictionary d to find the corresponding Python function to execute. The format of the dictionary is d = {key-1 : (Setup-1, Modify-1, Revert-1, Teardown-1), …, key-n : (Setup-n, Modify-n, Revert-n, Teardown-n) }. If the function-key matches, for

example, key-1 and the corresponding Alc-Dyn-Script-Action is “setup”, then the function specified as “Setup-1” will be executed. Setup and teardown functions are mandatory. Modify and revert functions are optional. If a modify function is defined, a corresponding revert function must also be defined. If no modify/revert function is required, the keyword “None” should be used instead.

- dyn.add_cli(string)

  This function is used to generate CLI output in the Python script. The use of dyn.add_cli (”””) allows the specification of strings spanning multiple lines, which drastically improves the readability of the script.

  A subset of all available CLI commands is currently enabled for dynamic data services. The command “tools dump service dynamic-services command-list” provides a complete overview of all available CLI nodes for dynamic data services. In the allowed nodes section, all CLI nodes are listed that can be navigated to and where attributes can be modified. The pass through nodes section shows CLI nodes that can be navigated to but no attribute changes are allowed. For example, it is not allowed to change the autonomous system of a router (configure router autonomous-system <autonomous-system>) because “configure router” is a “pass through node”. However, you can navigate to configure router, because you can add a static route: “/configure router static-route 0.0.0.0/0 next-hop 192.168.1.1” is part of the “allowed nodes”.

- dyn.select_free_id(“service-id”)

  This function is used to select a free service ID within the service ID range defined under dynamic-services context. An automatic assignment of the service id is one option, but it is also possible to provide the service id as one of the parameters in the “Alc-Dyn-Serv-Script-Params” list from RADIUS.

  The service-ID is a node-internal attribute. As such it is valid to let the node select the ID itself. However, in a network with multiple BNGs and a single customer service spanning two or more BNGs, a network administrator may actually prefer to use the same service-id for this customer service on all nodes for better visibility, which cannot be guaranteed if the automatic option is chosen. 5620 SAM is also using the service-ID as one attribute in addition to others to discover service-entities across the whole network. If SAM is in use for general management and service assurance, it is advised to manually specify the service-ID and not to use the automatic selection.

  In any case, the administrator needs to make a choice between the automatic ID assignment and the specific assignment for all dynamic data services, as a mix between both is not recommended.

  When the automatic assignment is chosen, there is no “binding/memory” of a service ID to a provisioned service, which means a service that may have service ID xyz initially may get another service ID the next time it comes up. In other words, as soon as a service is disconnected, the service ID is freed up for the next activated service.

- dyn.get_sap()

  This function returns the value of the evaluation of the “Alc-Dyn-Serv-SAP-Id” attribute as a string. Wildcards (“#”) in the Alc-Dyn-Serv-SAP-Id are replaced with the corresponding port/vlan information of the control channel SAP-ID. So if, for example, the “Alc-Dyn-Serv-SAP-Id” contains “#:#.1” and the control channel SAP ID is “1/1/5:100.100”, the resulting SAP for the data service would be “1/1/5:100.1”.

- dyn.get_circuit_id()

  This function returns a string which is equal to the Control Channel Circuit-ID (from the DHCP relay agent option 82 or PPP tags). This function may be useful, for example, to use the circuit id in the SAP description.

- dyn.get_remote_id()

  This function returns a string which is equal to the Control Channel Remote-ID (from the DHCP relay agent option 82 or PPP tags). This function may be useful, for example, to use the remote id in the SAP description.

In addition to the RADIUS dictionary, the node will also store service-related parameters in a service-specific dictionary. The information in the RADIUS messages or in the stored dictionary are used for the various functions as outlined in Table 43:

Table 43 **Function and Dictionary Relationship**

Function Name	Input	Returns
setup dynsvc(rd*)	**rd:** radius dictionary in the parameter list in Alc-Dyn-Serv-Script-Params. VSA Passed to setup function.	A dictionary that will be stored for the lifetime of the dynamic service (sd).
modify_dynsvc(rd,sd**)	**rd:** radius dictionary in the parameter list in Alc-Dyn-Serv-Script-Params. VSA passed to modify function. **sd:** previously stored dictionary of the setup/ previous modify functions.	Updated stored dictionary (sd)
revert_dynsvc(rd, sd)	**rd:** radius dictionary in the parameter list in Alc-Dyn-Serv-Script-Params. VSA passed to revert function. **sd:** previously stored dictionary of the setup/ previous modify function.	The function does not return (store) any information. The previously stored dictionary (sd) is kept.
teardown_dynsvc(sd)	**sd:** previously stored dictionary by the setup function or a previous modify function are passed to the teardown function.	The function does not return (store) any information. The stored dictionary (sd) is deleted.
(*) rd = radius dictionary		

Table 43 **Function and Dictionary Relationship (Continued)**

Function Name	Input	Returns
(**) sd = stored dictionary. sd is required for modifies, reverts and teardowns.		

Building Block: CLI Snippets

The necessary functional parts of a service configuration cannot typically be put into one large script (one single actionable function). This is best described with a small and simple example:

Imagine a single script where the setup action creates both the service instance and the SAP, and the teardown action removes the service instance and the SAP. For a service with just one SAP per service instance this may work fine, however, in a multi SAP service like a VPLS this will cause problems, especially during the service teardown action. This is because if multiple SAPs have been instantiated in a single service, the disconnect of just one SAP would trigger the teardown action which would try to remove the SAP (still ok) but then would try to remove the service instance. This action would fail as other SAPs still exist in the service. As such the script execution would fail.

It is therefore necessary to structure the whole required configuration into individual actionable pieces which are referenced by each other with specific reference-IDs. Those actionable pieces are called “snippets”.

Referenced snippets may or may not be executed depending on whether the functional instance exists already or not. As shown on the left of the picture below, the action to create a SAP references the creation of a service and then to the creation of a customer. For the very first business site to come up all three snippets will be executed. For any further business site to come up in the same service the script to create the SAP will be executed, the referenced service script and subsequently the customer script will not be executed again as those instances already exist. The same logic applies during the teardown action. Only when the last SAP in a service is removed is the service-instance itself removed, and potentially also the customer (unless it too is associated with other services).

Figure 209 **Hierarchy of Snippets**

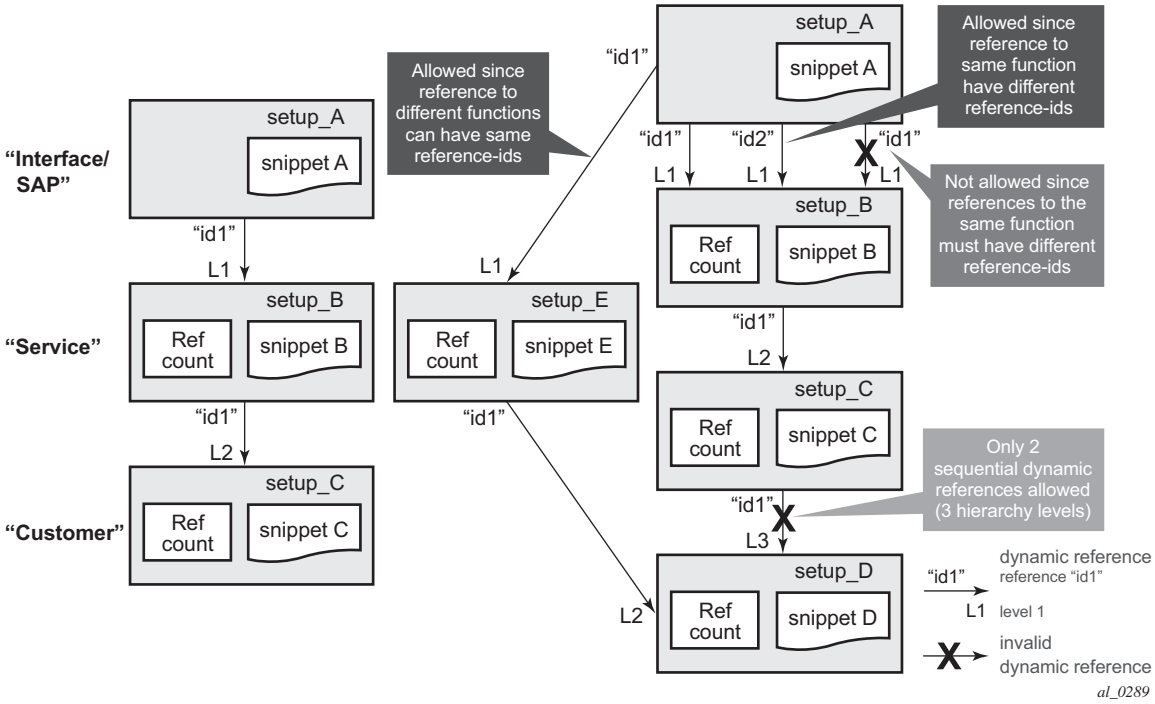

The implementation supports a three level hierarchy of snippets for high flexibility as shown in the picture. A reference to the fourth level as shown on the right side would result in an error.

Furthermore, snippets can be scaled "horizontally", so from one level multiple references to other snippets are possible. An example for that would be the creation of a SAP triggers the creation of a service as well as the creation of an Ethernet CFM association for that SAP.

Identifiers are needed for the referencing. The same identifier can be used on the "horizontal" level, but not on the vertical level between the same pair of snippets, also shown above.

Snippets are heavily used in the service examples in Bringing it all together where the logic and the referencing are described with real data.

Building Block: RADIUS Accounting

As dynamic data services are instantiated through RADIUS, it is also typically required to provide feedback to the RADIUS server for service establishment and teardown. This is achieved via RADIUS accounting records for the dynamic data channels in addition to the accounting messages for the PPP or DHCP control channel.

Up to two dedicated accounting destinations can be defined within the dynamic services policy. Thus, the accounting for the dynamic data services can be handled by an independent set of accounting servers (from the accounting for general ESM subscribers). But the same servers can also be used.

Each dynamic data service has its own accounting start/stop/interim messages based on a unique accounting session ID. In addition, the accounting packets contain a multi-session ID which is identical to the accounting session ID of the control channel and is therefore displayed in show commands as Acct-Session-ID-Ctrl as shown below.

```
A:BNG-1# show service dynamic-services saps summary
===============================================================================
Dynamic Services SAP's summary
===============================================================================
SAP                          Acct-Session-ID       Acct-Session-ID-Ctrl
-------------------------------------------------------------------------------
3/2/1:4.3                    D6E559000000B951668AEB D6E559000000B851668AEB
3/2/2:1.1                    D6E559000000C75166CFF4 D6E559000000C45166CFF4
3/2/2:1.2                    D6E559000000C85166CFF4 D6E559000000C45166CFF4
3/2/2:1.3                    D6E559000000C95166CFF4 D6E559000000C45166CFF4
3/2/2:1.4                    D6E559000000CA5166CFF4 D6E559000000C45166CFF4
-------------------------------------------------------------------------------
No. of SAP's: 5
===============================================================================
```

The Accounting Session ID (in the centre above) is the one for the dynamic data service itself, the one on the right is from the control-channel. The above example clearly shows that the last 4 dynamic services all belong to the same control channel, as they all have the same Acct-Session-ID-Ctrl.

If the accounting stats-type is set to “volume-time”, the interim and stop accounting messages will also contain counters for the data traffic through the service. With the accounting stats-type “time”, no counters are included, only session time is reported.

As a dynamic data service is functionally no different from a regular data service, traffic volumes can also be gathered by assigning accounting policies within the service for file-based XML accounting.

Bringing it all together

This section gives examples of all of the above parameters and will also cover show, log and debug information.

In the given example, a single user in the database has four different associated data services. Not only are the data service types all different, but also other aspects of the parameter set, this has an effect on how the data is entered in the RADIUS VSAs and how the Python script is constructed. More detail is given below. The different models for specifying parameters are presented to show the flexibility. An operator typically chooses a single model and uses that for all its services.

As all of the information for these four services will potentially be sent in one RADIUS message, the VSAs need to be tagged so that the BNG can link the appropriate VSAs to each other and differentiate the services. For better visibility, the different sections in the RADIUS users file are displayed with bold black and dark grey text.

The freeradius users file format is used for this example.

```
1.    "subscriber12@domain2.com" Cleartext-Password := "ALU"
2.        Alc-Subsc-ID-Str := "pppoe-user12",
3.        Framed-IP-Address = 10.2.1.200,
4         Alc-Dyn-Serv-SAP-Id:1  = "#:#.1",
5.        Alc-Dyn-Serv-Script-Params:1 = "business_epipe={'t':('EPipe-
6.            CustomerName','CustomerName-Circuit-1','3','3','64496',
7.            '192.0.2.5','192.0.2.1','3333')}",
8.        Alc-Dyn-Serv-Policy:1 =  "dynamic-services-2",
9.        Alc-Dyn-Serv-SAP-Id:2 += "#:#.2",
10.       Alc-Dyn-Serv-Script-Params:2 += "business_vprn={'t':('9999',
11.           'VPRN-CustomerName','64497','100000','CustomerName-Circuit-
12.           1','172.16.10.1/30','3','1','3','2','172.16.100.0/24',
13.           '172.16.10.2','100')}",
14.       Alc-Dyn-Serv-Acct-Interim-Ivl-1:2 += "600",
15.       Alc-Dyn-Serv-Acct-Interim-Ivl-2:2 += "0",
16.       Alc-Dyn-Serv-Policy:2  +=  "dynamic-services-2",
17.       Alc-Dyn-Serv-SAP-Id:3  += "#:#.3",
18.       Alc-Dyn-Serv-Script-Params:3 += "business_vpls={'inst':
19.           'VPLS-CustomerName','if_name':'CustomerName-Circuit-1','ing_qos':'3',
20.           'egr_qos':'3','imp_comm_val':'10000','exp_comm_val':'10000',
21.           'rt':'64498','rd':'64498'}",
22.       Alc-Dyn-Serv-Policy:3 +=  "dynamic-services-2",
23.       Alc-Dyn-Serv-Acct-Interim-Ivl-1:3 += "0",
24.       Alc-Dyn-Serv-Acct-Interim-Ivl-2:3 += "0",
25.       Alc-Dyn-Serv-Acct-Stats-Type-1:3 += off,
26.       Alc-Dyn-Serv-Acct-Stats-Type-2:3 += off,
27.       Alc-Dyn-Serv-SAP-Id:4  += "#:#.4",
28.       Alc-Dyn-Serv-Script-Params:4 += "business_ies={'t':
29.           ('IES-CustomerName','CustomerName-Circuit-1','172.16.11.1/30',
30.           '2001:db8:5100:1000::1/64','5','1','1','6','2','2','5','25',
31.           'cfm-Mep-to-CPE','100',",
32.       Alc-Dyn-Serv-Script-Params:4 += "[{'to':'172.16.110.0/24',
33,           'n-h':'172.16.11.2'},{'to':'2001:db8:bbbb::/56',
34.           'n-h':'2001:db8:5100:1000::2'}])}",
```

```
35.       Alc-Dyn-Serv-Policy:4 +=  "dynamic-services-2",
36.       Alc-Dyn-Serv-Acct-Interim-Ivl-1:4 += "600",
37.       Alc-Dyn-Serv-Acct-Interim-Ivl-2:4 += "0",
38.       Alc-Dyn-Serv-Acct-Stats-Type-1:4 += "3",
39.       Alc-Dyn-Serv-Acct-Stats-Type-2:4 += "2",
```

The first section (lines 1 — 3) shows a minimal parameter set for the (PPP) control channel. As the focus of this example is on the dynamic data services, all default parameters will be used for the control-session which are defined under the msap.

The second section (lines 4 — 8, attributes with tag ":1") shows a possible parameter set for an Epipe service. Only the absolutely minimum set of VSAs is used (see Dynamic Service Attribute List for Setup, Modify and Teardown). Furthermore, all service parameters are listed without keywords in a pre-defined order. No service ID number is specified in "Alc-Dyn-Serv-Script-Params", hence the Python script should dynamically select the next free ID.

The third section (lines 9 — 16, attributes with tag ":2") shows a possible parameter set for a VPRN service. A few more VSAs are defined, thus some of the default parameters in the dynamic service policy are overwritten for this service. The first entry in the "Alc-Dyn-Serv-Script-Params" attribute specifies the Service-ID number for this service, so the Python script should not select a service ID automatically. Furthermore, static-routing information towards the CPE is added as normal attributes at the end of the list.

The fourth section (line 17 — 26, attributes with tag ":3") shows a possible parameter set for a VPLS service. Notice the difference with the first two services in the "Alc-Dyn-Serv-Script-Params" part: now all parameters are given their specific keyword. As such, the sequence of those parameters is not important. The effect on the Python script is shown further down.

The fifth section (lines 27 — 39, attributes with tag ":4") finally shows a possible parameter set for an IES service. All of the required parameters for this service do not fit into a single "Alc-Dyn-Serv-Script-Params" attribute anymore (limited to 247 bytes). As is shown, multiple VSAs can be "concatenated" by simply splitting the attributes. It is important that the order in which the different "Alc-Dyn-Serv-Script-Params" attributes with the same tag is received can be guaranteed. Furthermore the second appearance of this VSA shows a different way of provisioning static-routing information towards the CPE.

To better understand the details it is necessary to take a closer look into the active Python script. The first important part is the section with the dynamic actions.

```
-snip-

d = {
"vprn": (setup_vprn, None, None, teardown_vprn),
"ies": (setup_ies, None, None, teardown_ies),
"vpls": (setup_vpls, None, None, teardown_vpls),
```

```
"epipe": (setup_epipe, None, None, teardown_epipe),
"ethcfm" : (setup_ethcfm_domain, None, None, teardown_ethcfm_domain),
"business_vprn" : (setup_business_vprn, None, None, teardown_business_vprn),
"business_ies" : (setup_business_ies, None, None, teardown_business_ies),
"business_vpls" : (setup_business_vpls, None, None, teardown_business_vpls),
"business_epipe" : (setup_business_epipe, modify_business_epipe,
                    revert_business_epipe, teardown_business_epipe)}
dyn.action(d)
```

The function-key string specified at the start of the “Alc-Dyn-Serv-Script-Params” (for example Alc-Dyn-Serv-Script-Params:1 = “business_epipe={…}”) has a 1:1 mapping with the keys of the dictionary “d” in the highlighted section of the above sample (for example d = { …, "business_epipe" : (…)}. For services, different values for setup, modify, revert and teardown are given which point to other sections in the Python script (see below). Setup and teardown functions are mandatory, whereas modify and revert functions are optional.

In the unbolded text of the previous example, there are other actions defined that are not contained in the RADIUS attributes (for example d = { "vprn": (…) , …). Those actions are referenced by the four main functions.

In the next part, there is more detail presented in each service example and maps it to the corresponding Python function.

It is advisable to read through all examples, as only the deltas between each service are explicitly explained.

Example 1 – Epipe service

```
# copy of the RADIUS attributes from above
-snip-
        Alc-Dyn-Serv-SAP-Id:1  = "#:#.1",
        Alc-Dyn-Serv-Script-Params:1 = "business_epipe={'t':
            ('EPipe-CustomerName','CustomerName-Circuit-1','3','3','64496'
             ,'192.0.2.5','192.0.2.1','3333')}",
        Alc-Dyn-Serv-Policy:1 =  "dynamic-services-2",
-snip-

# Python-part
d = {
-snip-
"business_epipe" : (setup_business_epipe, modify_business_epipe,
                    revert_business_epipe, teardown_business_epipe)
dyn.action(d)
-snip-
def setup_business_epipe(d):
    keys = ('inst', 'if_name', 'ing_qos', 'egr_qos', 'as', 'remote_ip',
           'local_ip', 'glb_svc_id')
    d = dict(zip(keys, d['t']))
    ref_d = dyn.reference("epipe", d['inst'], d)
    d['svc_id'] = ref_d['svc_id']
    d['sap_id'] = dyn.get_sap()
    dyn.add_cli("""
```

```
configure
  service
    epipe %(svc_id)s
      sap %(sap_id)s create
        description "%(if_name)s"
        ingress
          qos %(ing_qos)s
        exit
        egress
          qos %(egr_qos)s
        exit
      exit
      spoke-sdp-fec %(svc_id)s fec 129 aii-type 2 create
        pw-template-bind 2
        saii-type2 %(as)s:%(local_ip)s:%(glb_svc_id)s
        taii-type2 %(as)s:%(remote_ip)s:%(glb_svc_id)s
        no shutdown
      exit
    exit
  exit
exit
""" % d)
    return d

def setup_epipe(d):
    d['svc_id'] = dyn.select_free_id("service-id")
    dyn.add_cli("""
configure
  service
    epipe %(svc_id)s customer 1 create
      service-name "%(inst)s"
      description "%(inst)s"
      no shutdown
    exit
  exit
exit
""" % d)
    return {'svc_id':d['svc_id']}

def teardown_epipe(d):
    dyn.add_cli("""
configure
  service
    epipe %(svc_id)s
      shutdown
    exit
    no epipe %(svc_id)s
  exit
exit
""" % d)

def teardown_business_epipe(d):
    dyn.add_cli("""
configure
  service
    epipe %(svc_id)s
      sap %(sap_id)s
        shutdown
      exit
```

```
      spoke-sdp-fec %(svc_id)s
        shutdown
      exit
      no sap %(sap_id)s
      no spoke-sdp-fec %(svc_id)s
    exit
  exit
exit
""" % d)
-snip-
```

Based on the dictionary specified in the dyn.action(d) call, the function definition “setup_business_epipe” in the Python script corresponds with the function that will be called if the function-key “business-epipe” is specified in the “Alc-Dyn-Serv-Script-Params” attribute as dictionary name and if a setup action is required. The dictionary containing the parameters in the RADIUS VSA “Alc-Dyn-Serv-Script-Params” has a single key-value pair, with the parameters stored in a tuple. The individual parameters cannot be identified with a keyword hence the order in which they are specified in the RADIUS VSA should match the order in which they are extracted in the Python script. The first two lines in this part of the script extract the parameters out of the array “t” and link them to unique keywords, which are used for the rest of the script.

The parameter “inst” is important in this logic, as it defines whether access circuits belong to the same service-instance or different instances (the RADIUS VSAs for two SAPs belonging to the same service therefore need to have the same “inst” value). If you look at the CLI of the “setup_business_epipe” function, you can see that it creates the SAP and all related attributes, but not the service itself. It is the “ref_d = dyn.reference("epipe", d['inst'], d)” that references a part in the script to create the actual service-instance. The referenced function is found by using the first parameter in the dyn.reference call (“epipe”) as a function-key lookup in the dictionary specified in the dyn.action(d) and finding the corresponding setup function: d = { …, "epipe" : (setup_epipe, …), …}. The second parameter (“d['inst']”) is used as unique identification of the service instance. The last parameter (“d”) is a dictionary with parameters that can be used by the references function. When the first customer endpoint with a new “inst” name comes up, the service itself gets created.

By looking at “def setup_epipe(d):” the first line “d['svc_id'] = dyn.select_free_id("service-id")” of the script automatically picks a free service-id out of the range defined in the dynamic service policy, as no service ID was provided in the RADIUS parameters. The rest of this function creates the service instance. Service attributes that were provided by RADIUS and are placed in a service specific dictionary are available to this function via the third parameter in the dyn.reference call. The newly generated service ID is returned to the calling script by the “return {'svc_id':d['svc_id']}” command at the end of the function. The service specific dictionary (as explained in the Python Script Building Block) is updated with the appropriate information.

Back to “def setup_business_epipe(d):”, the service ID together with the SAP ID and the parameters from the Alc-Dyn-Serv-Script-Params VSA are used to create the appropriate CLI code for the SAP and the SDP within the service.

Similar to the setup, there is also a teardown part for both service and SAP. The teardown function is called either through the termination of the control-channel, through a COA with Alc-Dyn-Script-Action = teardown or through a disconnect message. The CLI for the teardown script must be written in the correct sequence as applied by the SR OS CLI logic so that SAP(s) and service(s) are removed in the correct order.

Example 2 – VPRN service

```
RADIUS-part from above
-snip-
        Alc-Dyn-Serv-SAP-Id:2 += "#:#.2",
        Alc-Dyn-Serv-Script-Params:2 += "business_vprn={'t':
             ('9999','VPRN-CustomerName','64497','100000',
              'CustomerName-Circuit-1','172.16.10.1/30','3','1','3','2',
              '172.16.100.0/24','172.16.10.2','100')}",
        Alc-Dyn-Serv-Acct-Interim-Ivl-1:2 += "600",
        Alc-Dyn-Serv-Acct-Interim-Ivl-2:2 += "0",
        Alc-Dyn-Serv-Policy:2  +=  "dynamic-services-2",
-snip-

Python-part
d = {
-snip-
"business_vprn" : (setup_business_vprn, None, None, teardown_business_vprn)
dyn.action(d)
-snip-
def setup_business_vprn(d):
    keys = ('svc_id', 'inst', 'as_id', 'comm_id', 'if_name', 'ipv4_address',
            'ing_qos', 'ing_ip_filter', 'egr_qos', 'egr_ip_filter', 'lan_pfx',
            'nxt_hop', 'metric')
    d = dict(zip(keys, d['t']))
    ref_d = dyn.reference("vprn", d['inst'], d)
    d['sap_id'] = dyn.get_sap()
    dyn.add_cli("""
configure
  service
    vprn %(svc_id)s
      interface "%(if_name)s" create
        address %(ipv4_address)s
        urpf-check mode strict
        sap %(sap_id)s create
          ingress
            qos %(ing_qos)s
            filter ip %(ing_ip_filter)s
          exit
          egress
            qos %(egr_qos)s
            filter ip %(egr_ip_filter)s
          exit
        exit
      exit
```

```
    exit
  exit
  router
    static-route %(lan_pfx)s next-hop %(nxt_hop)s metric %(metric)s
  exit
exit
""" % d)
    return d

def setup_vprn(d):
    dyn.add_cli("""
configure
  service
    vprn %(svc_id)s customer 1 create
      service-name "%(inst)s"
      description "%(inst)s"
      autonomous-system %(as_id)s
      route-distinguisher %(as_id)s:%(comm_id)s
      auto-bind mpls
      vrf-target target:%(as_id)s:%(comm_id)s
      no shutdown
    exit
  exit
exit
""" % d)
    return {'svc_id':d['svc_id']}

def teardown_vprn(d):
    dyn.add_cli("""
configure
  service
    vprn %(svc_id)s
      shutdown
    exit
    no vprn %(svc_id)s
  exit
exit
""" % d)

def teardown_business_vprn(d):
    dyn.add_cli("""
configure
   router
     no static-route %(lan_pfx)s next-hop %(nxt_hop)s
   exit
   service
     vprn %(svc_id)s
       interface "%(if_name)s"
         sap %(sap_id)s
           shutdown
         exit
         no sap %(sap_id)s
         shutdown
       exit
       no interface "%(if_name)s"
     exit
   exit
exit
""" % d)
```

```
-snip-
```

In this example of a VPRN service two additional RADIUS VSAs are used to overwrite the accounting interim update intervals for the two RADIUS Accounting servers that are specified in the dynamic services policy. The Stats-Type configuration (time or volume-time) is obtained from the dynamic services policy as no RADIUS VSA is provided for that.

The beginning of the “setup_business_vprn” definition is identical to the earlier Epipe service example. This time a service identifier is provided as part of the parameter list. The referenced function to create the VPRN service (def setup_vprn) does not need the line to auto-generate the service ID.

At the end of the setup-procedure there is a basic example to add static-route information in case they are needed for PE-CE communication. Later on, in the IES service example, a more flexible alternative is shown.

Example 3 – VPLS service

```
RADIUS-part from above
-snip-
        Alc-Dyn-Serv-SAP-Id:3  += "#:#.3",
        Alc-Dyn-Serv-Script-Params:3 += "business_vpls={'inst':
            'VPLS-CustomerName','if_name':'CustomerName-Circuit-1','ing_qos':'3',
            'egr_qos':'3','imp_comm_val':'10000','exp_comm_val':'10000',
            'rt':'64498','rd':'64498'}",
     Alc-Dyn-Serv-Policy:3 +=  "dynamic-services-2",
        Alc-Dyn-Serv-Acct-Interim-Ivl-1:3 += "0",
        Alc-Dyn-Serv-Acct-Interim-Ivl-2:3 += "0",
        Alc-Dyn-Serv-Acct-Stats-Type-1:3 += off,
        Alc-Dyn-Serv-Acct-Stats-Type-2:3 += off,
-snip-

Python-part
d = {
-snip-
"business_vpls" : (setup_business_vpls, None, None, teardown_business_vpls)
-snip-
def setup_business_vpls(d):
    ref_d = dyn.reference("vpls", d['inst'], d)
    d['svc_id'] = ref_d['svc_id']
    d['sap_id'] = dyn.get_sap()
    dyn.add_cli("""
configure
   service
    vpls %(svc_id)s
      sap %(sap_id)s create
        description "%(if_name)s"
        ingress
          qos %(ing_qos)s
        exit
        egress
          qos %(egr_qos)s
        exit
```

```
        collect-stats
        accounting-policy 10
      exit
    exit
  exit
exit
""" % d)
    return d

def setup_vpls(d):
    d['svc_id'] = dyn.select_free_id("service-id")
    dyn.add_cli("""
configure
  service
    vpls %(svc_id)s customer 1 create
      service-name "%(inst)s"
      description "%(inst)s"
      bgp
        route-distinguisher %(rd)s:%(exp_comm_val)s
        route-target export target:%(rt)s:%(exp_comm_val)s
                      import target:%(rt)s:%(imp_comm_val)s
        pw-template-binding 1
        exit
      exit
      bgp-ad
        vpls-id %(rt)s:%(exp_comm_val)s
        no shutdown
      exit
      no shutdown
    exit
  exit
exit
""" % d)
    return {'svc_id':d['svc_id']}

def teardown_vpls(d):
    dyn.add_cli("""
configure
  service
    vpls %(svc_id)s
      shutdown
      bgp-ad
        shutdown
      exit
      no bgp-ad
      bgp
        no pw-template-binding 1
      exit
    exit
    no vpls %(svc_id)s
  exit
exit
""" % d)

def teardown_business_vpls(d):
    dyn.add_cli("""
configure
  service
    vpls %(svc_id)s
```

```
      sap %(sap_id)s
        shutdown
      exit
      no sap %(sap_id)s
    exit
  exit
exit
""" % d)
-snip-
```

In the VPLS example the “Alc-Dyn-Serv-Acct-Stats-Type” is set to “off” for both RADIUS accounting destinations, meaning RADIUS accounting is switched off, even if it is enabled in the dynamic data services policy. In the script you can see that this service uses XML-accounting on the SAP instead (“collect-stats” and “accounting-policy 10”).

The dictionary containing the parameters in the RADIUS VSA “Alc-Dyn-Serv-Script-Params” has a key-value pair for each parameter. In the Python script the individual parameters can be identified immediately with the dictionary key. The order in which they are specified in the RADIUS VSA does not have to be strictly defined. The drawback of this approach is that the length of the parameter VSA increases. A single parameter VSA is limited to a length of 246 bytes and the total length of all parameter VSAs for a single service is limited to 1000 bytes.

Example 4 – IES service

```
RADIUS-part from above
-snip-
        Alc-Dyn-Serv-SAP-Id:4  += "#:#.4",
        Alc-Dyn-Serv-Script-Params:4 += "business_ies={'t':
            ('IES-CustomerName','CustomerName-Circuit-1','172.16.11.1/30',
             '2001:db8:5100:1000::1/64','5','1','1','6','2','2','5','25',
             'cfm-Mep-to-CPE','100',",
        Alc-Dyn-Serv-Script-Params:4 += "[{'to':'172.16.110.0/24',
            'n-h':'172.16.11.2'},{'to':'2001:db8:bbbb::/56',
            'n-h':'2001:db8:5100:1000::2'}])}",
        Alc-Dyn-Serv-Policy:4 +=  "dynamic-services-2",
        Alc-Dyn-Serv-Acct-Interim-Ivl-1:4 += "600",
        Alc-Dyn-Serv-Acct-Interim-Ivl-2:4 += "0",
        Alc-Dyn-Serv-Acct-Stats-Type-1:4 += "3",
        Alc-Dyn-Serv-Acct-Stats-Type-2:4 += "2",
-snip-

Python-part
d = {
-snip-
"business_ies" : (setup_business_ies, None, None, teardown_business_ies)
-snip-
def setup_business_ies(d):
    keys = ('inst', 'if_name', 'ipv4_address', 'ipv6_address', 'ing_qos',
            'ing_ip_filter', 'ing_ipv6_filter',  'egr_qos', 'egr_ip_filter',
            'egr_ipv6_filter', 'ing_bw', 'egr_bw', 'cfm_assoc_id', 'metric',
            'routes')
```

```
    d = dict(zip(keys, d['t']))
    ref_d = dyn.reference("ies", d['inst'], d)
    d['svc_id'] = ref_d['svc_id']
    d['sap_id'] = dyn.get_sap()
    d['cfm_domain'] = 1
    ref_d_cfm = dyn.reference("ethcfm", str(d['cfm_domain']), d)
    dyn.add_cli("""
configure
  eth-cfm
    domain %(cfm_domain)s
      association %(svc_id)s format string name "%(cfm_assoc_id)s"
        bridge-identifier %(svc_id)s
        exit
        ccm-interval 1
        remote-mepid 2
      exit
    exit
  exit
  service
    ies %(svc_id)s
      interface "%(if_name)s" create
        address %(ipv4_address)s
        urpf-check mode strict
        cflowd interface both
        ipv6
          address %(ipv6_address)s
          urpf-check mode strict
        exit
        sap %(sap_id)s create
           description "%(if_name)s"
          ingress
            scheduler-policy "Business Services"
            scheduler-override
               scheduler "root-t1" create
                  rate %(ing_bw)s000
               exit
            exit
            qos %(ing_qos)s
            filter ip %(ing_ip_filter)s
            filter ipv6 %(ing_ipv6_filter)s
          exit
          egress
            qos %(egr_qos)s
            filter ip %(egr_ip_filter)s
            filter ipv6 %(egr_ip_filter)s
            agg-rate-limit %(egr_bw)s000 queue-frame-based-accounting
          exit
          collect-stats
          accounting-policy 10
          eth-cfm
            mep 1 domain %(cfm_domain)s association %(svc_id)s direction down
              ccm-enable
              no shutdown
            exit
          exit
        exit
        urpf-check
        exit
      exit
```

```
    exit
  exit
  router
""" % d)
    for route in d['routes']:
        dyn.add_cli("""
        static-route %s next-hop %s metric %s tag 80
"""% (route["to"], route["n-h"], d['metric']))
    dyn.add_cli("""
    exit
exit
""" % d)
    return d

def setup_ies(d):
    d['svc_id'] = dyn.select_free_id("service-id")
    dyn.add_cli("""
configure
  service
    ies %(svc_id)s customer 1 create
      service-name "%(inst)s"
      description "%(inst)s"
      no shutdown
    exit
  exit
exit
""" % d)
    return {'svc_id':d['svc_id']}

def setup_ethcfm_domain(d):
    dyn.add_cli("""
configure
  eth-cfm
    domain %(cfm_domain)s format none level 1
    exit
  exit
exit
""" % d)
    return {'cfm_domain':d['cfm_domain']}

def teardown_ethcfm_domain(d):
    dyn.add_cli("""
configure
  eth-cfm
    no domain %(cfm_domain)s
  exit
exit
""" % d)

def teardown_ies(d):
    dyn.add_cli("""
configure
  service
    ies %(svc_id)s
      shutdown
    exit
    no ies %(svc_id)s
  exit
exit
```

```
""" % d)

def teardown_business_ies(d):
    dyn.add_cli("""
configure
  router
""")
    for route in d['routes']:
        dyn.add_cli("""
        no static-route %s next-hop %s
"""% (route["to"], route["n-h"]))
    dyn.add_cli("""
   exit
exit
""")
    dyn.add_cli("""
configure
  service
    ies %(svc_id)s
      interface "%(if_name)s"
        sap %(sap_id)s
          shutdown
          eth-cfm
            mep 1 domain %(cfm_domain)s association %(svc_id)s
              shutdown
            exit
            no mep 1 domain %(cfm_domain)s association %(svc_id)s
          exit
        exit
        no sap %(sap_id)s
        shutdown
      exit
      no interface "%(if_name)s"
    exit
  exit
  eth-cfm
    domain %(cfm_domain)s
      association %(svc_id)s
        no bridge-identifier %(svc_id)s
      exit
      no association %(svc_id)s
    exit
  exit
exit
""" % d)
-snip-
```

The IES example has the most attributes. The maximum length of a tagged RADIUS VSA is 246 bytes. If the amount of data is too big to fit into one attribute, simply add a second or third one in the syntax shown above in the RADIUS part. There is no need to separate the attributes exactly at 246 bytes; it can be cut at any position in the list (preferably between two attributes for better readability). Note also that all the parameter VSAs that belong to the same service should have the same tag (“:4” in this example).

In case of multiple parameter VSAs, the order in which they are specified is important and must be guaranteed as the concatenation of all the attributes must result in a Python dictionary in the form: “dictionary-name = {…}”. The Python script is not aware that multiple attributes were used.

Another difference to the previous examples is that there is not only a reference to the function for the service creation, but also a similar reference to a function for Ethernet Connectivity Fault Management (CFM). Considering that you may want to put all of the Eth-CFM endpoints under the same domain within unique associations, the Eth-CFM domain needs to be created first and torn down as last.

Finally, a different way to provide static-route information is shown at the end of the “setup_business_ies” definition (starting with “for route in d['routes']:”). Also note the difference in how this information is implemented at the end of the “Alc-Dyn-Serv-Script-Params” list. The static routes themselves are defined as a dictionary and thus as many routes as required can be added with this method. Compare this to the VPRN example where a more basic mechanism was used.

As outlined before, dynamic data services can be triggered during the Access-Accept for the control channel but also through a CoA to the control channel Accounting Session ID.

Example 5 – modify an Epipe service using CoA

So far the focus was on service establishment and teardown. It is also possible to change a running dynamic data service using the “modify” function. This will be explained with the previously configured Epipe service.

```
RADIUS attributes in the COA message
   Acct-Session-Id = D6E559000000BD5166BF34 #
   Alc-Dyn-Serv-SAP-Id:1  = "#:#.1",
   Alc-Dyn-Serv-Script-Params:1 = "business_epipe={'ing_qos':'4','egr_qos':'4'}",
   Alc-Dyn-Serv-Script-Action:1 = modify,
   Alc-Dyn-Serv-Policy:1 =  "dynamic-services-2",


Python-part
d = {
-snip-
"business_epipe" : (setup_business_epipe, modify_business_epipe, revert_business_epi
pe, teardown_business_epipe)}
dyn.action(d)
-snip-
def modify_business_epipe(d, sd):
    sd['ing_qos'] = d['ing_qos']
    sd['egr_qos'] = d['egr_qos']
    dyn.add_cli("""
configure
   service
      epipe %(svc_id)s
         sap %(sap_id)s
           ingress
```

```
                qos %(ing_qos)s
              exit
              egress
                qos %(egr_qos)s
              exit
            exit
          exit
        exit
exit
"""% sd)
    return sd

def revert_business_epipe(d, sd):
    dyn.add_cli("""
configure
   service
     epipe %(svc_id)s
        sap %(sap_id)s
          ingress
            qos %(ing_qos)s
          exit
          egress
            qos %(egr_qos)s
          exit
        exit
      exit
    exit
exit
""" % sd)
-snip-
```

Through the function-key in the parameter list (Alc-Dyn-Serv-Script-Params:1 = "business_epipe= …) and the action attribute of “modify” (Alc-Dyn-Serv-Script-Action:1 = modify), the script will identify the relevant routine to be invoked for the modification (modify_business_epipe). If a modify function is defined, there must also be a definition for a revert function. A revert function cannot be initiated from RADIUS, but it is automatically executed to restore the initial configuration in case the modify script execution fails.

A modify action for an existing service is triggered with a CoA message. For this CoA, either the Accounting Session ID (ASID) of the control channel or the Accounting Session ID of the dynamic data channel can be used. In case the ASID of the control channel is used, the “Alc-Dyn-Serv-SAP-Id” can contain wildcards, as the appropriate port and VLAN information will be taken from the control channel. If the ASID of the dynamic data channel itself is used, the “Alc-Dyn-Serv-SAP-Id” needs to be fully specified, without wildcards. Otherwise the script execution will fail.

For a modify action, the “Alc-Dyn-Serv-Script-Params” only contains the parameters to be changed and does not need any further service identifying information. The service is identified based on the ASID and the “Alc-Dyn-Serv-SAP-Id”. Parameters which have been previously received by the setup or an earlier modify function are available in the stored dictionary (sd). Those are combined with the dictionary in the RADIUS message (d). Service modifications which relate to subsequent modifications, or for the service teardown, need to be updated in the stored dictionary so that they can be used in those later actions. This is achieved by the “return sd” command.

As with “manual” provisioned services, the new QoS settings from our example take effect immediately.

A dynamic data service can also be disconnected using a RADIUS Disconnect Message containing the Accounting Session ID of the dynamic data service, or indirectly via a RADIUS Disconnect Message containing the Accounting Session ID of the control channel which would result in a teardown of all associated dynamic data services.

Debugging

It is obvious that the Python scripts need extensive testing in the lab before they are deployed in the field. This testing may require a number of iterations: write the script, testing, verification, improvement and testing again. Every time there is a change in the Python script the node needs to reload the script. This is achieved by a **shutdown** and **no shutdown** of the active script using the command:

configure aaa radius-script-policy <script-policy-name> <primary/secondary> shutdown

configure aaa radius-script-policy <script-policy-name> <primary/secondary> no shutdown

Testing the script may result in some problems if certain aspects may not work as expected (see also debug functions later in this section). It can be that a dynamically created service cannot be removed properly because the teardown script contains errors and the whole service, or fragments of that service, may still exist on the node.

Dynamic data services cannot be edited in normal CLI mode as it may potentially make a later removal of that service through the script impossible. For troubleshooting there is a procedure to manipulate those services during the testing phase, thus avoiding the need to reboot the box to clear the state. The **enable-dynamic-services-config** command allows for the editing dynamic services just like normal services. As this is an action that should only be executed by authorized personnel, the activation of this command is protected by the use of a password, defined under **configure system security password dynsvc-password**.

The **show users** command has been extended to visualize the respective mode ('D' indicates a user is in dynamic service edit mode). A user in dynamic services edit mode cannot modify regular services.

no enable-dynamic-services-config returns the user to normal mode.

To support the creation and the troubleshooting during the test phase the SR OS debug functions have been extended extensively to allow for a detailed review of what is happening in the script and on the CLI.

```
debug dynamic-services
debug dynamic-services scripts
debug dynamic-services scripts event
debug dynamic-services scripts event cli
debug dynamic-services scripts event errors
debug dynamic-services scripts event executed-cmd
debug dynamic-services scripts event state-change
debug dynamic-services scripts event warnings
debug dynamic-services scripts instance
debug dynamic-services scripts instance event
debug dynamic-services scripts instance event cli
debug dynamic-services scripts instance event errors
debug dynamic-services scripts instance event executed-cmd
debug dynamic-services scripts instance event state-change
debug dynamic-services scripts instance event warnings
debug dynamic-services scripts script
debug dynamic-services scripts script event
debug dynamic-services scripts script event cli
debug dynamic-services scripts script event errors
debug dynamic-services scripts script event executed-cmd
debug dynamic-services scripts script event state-change
debug dynamic-services scripts script event warnings
```

It is advised to enable all debug options when starting and then remove more and more debugs options as the script becomes more complete and stable. The debug output gives clear indications about errors in the script or its execution in case something goes wrong.

An additional aid is the use of “print” commands in the Python script itself for certain attributes during the execution of the script. The print output will appear in the debug log. “Print” commands in the Python script should only be used during the testing phase and not in the normal operations mode.

The following command allows the execution of a dynamic services Python script without the need for RADIUS interaction:

tools perform service dynamic-services evaluate-script sap <sap-id> control-session <acct-session-id> action <script-action> [dynsvc-policy <name>]

show service dynamic-services script statistics provides general statistics about script execution.

show service dynamic-services script snippets displays the individual service configuration parts and allows to check if all “snippets” are actually referenced (the counter will increment/decrement with every function call).

In the case of a failed script action a SAP may not be deleted properly and it remains in the configuration as an “orphaned” object.

An orphaned object no longer has any references, which can be seen using **show service dynamic-services root-objects** where the snippet name and snippet instance is set to “N/A”.

Complete setup flow example

To finalize the section about the interaction between RADIUS and the Python script, the complete setup flow for the Epipe example is shown using extracts from the debug output (any missing sequence numbers in the flow below are simple acknowledge messages from RADIUS and are left out to focus on the important information). The debug settings to be used for this output are the following.

```
*A:BNG-1# show debug
debug
    router "Base"
        radius
            packet-type authentication accounting coa
            detail-level medium
        exit
    exit
    router "management"
        radius
            packet-type authentication accounting coa
            detail-level medium
        exit
    exit
    dynamic-services
        scripts
            event
                cli
```

```
            exit
            instance "dynamic-services-1"
                event
                    cli
                exit
            exit
        exit
    exit
exit
```

The first sequence in the flow is the Access-Request to the RADIUS server for the control channel. The information provided is that configured as part of the regular ESM configuration.

```
9 2013/04/12 20:47:23.73 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
Access-Request(1) 172.31.1.2:1812 id 70 len 206 vrid 1 pol authentication-2
USER NAME [1] 24 subscriber12@domain2.com
NAS IP ADDRESS [4] 4 192.0.2.1
SERVICE TYPE [6] 4 Framed(2)
FRAMED PROTOCOL [7] 4 PPP(1)
CHAP PASSWORD [3] 17 1 0xd4b73e0a17c0ad7f03c19bc1db5c291d
CHAP CHALLENGE [60] 41 0x620fa5f8be193d2066f6abad96c7de2df03986e3421f9733220d9520137
b0bf40b30edc9c92bea30a2
VSA [26] 29 DSL(3561)
AGENT CIRCUIT ID [1] 13 circuit-id-12
AGENT REMOTE ID [2] 12 remote-id-12
NAS PORT ID [87] 11 3/2/2:1.100
CALLING STATION ID [31] 17 00:00:64:01:02:03
NAS IDENTIFIER [32] 5 BNG-1
NAS PORT TYPE [61] 4 PPPoEoQinQ(34)
```

If the subscriber can be authenticated and authorized, RADIUS responds with an Access-Accept containing attributes for both the control channel and the dynamic data service.

```
10 2013/04/12 20:47:23.73 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 70 len 211 from 172.31.1.2:1812 vrid 1 pol authentication-2
    VSA [26] 14 Alcatel(6527)
      SUBSC ID STR [11] 12 pppoe-user12
    FRAMED IP ADDRESS [8] 4 10.2.1.200
    VSA [26] 8 Alcatel(6527)
      DYN SERV SAP ID [164] 6 1 #:#.1
    VSA [26] 118 Alcatel(6527)
      DYN SERV SCRIPT PARAMS [165] 116 1 business_epipe={'t':('EPipe-
CustomerName','CustomerName-Circuit-
1','3','3','64496','192.0.2.5','192.0.2.1','3333')}
    VSA [26] 21 Alcatel(6527)
      DYN SERV POLICY [167] 19 1 dynamic-services-2
```

The existence of the Dyn Serv VSAs in the response triggers the BNG to start the execution of the Python script, but first the control channel session is completely established and an accounting start message is send to RADIUS. This is a standard accounting message for ESM subscribers.

```
11 2013/04/12 20:47:23.75 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Accounting-Request(4) 172.31.1.2:1813 id 108 len 191 vrid 1 pol accounting-2
    STATUS TYPE [40] 4 Start(1)
    NAS IP ADDRESS [4] 4 192.0.2.1
    SERVICE TYPE [6] 4 Framed(2)
    FRAMED PROTOCOL [7] 4 PPP(1)
    FRAMED IP ADDRESS [8] 4 10.2.1.200
    FRAMED IP NETMASK [9] 4 255.255.255.255
    NAS IDENTIFIER [32] 5 BNG-1
    SESSION ID [44] 22 D6E559000000D2516872DB
    MULTI SESSION ID [50] 22 D6E559000000D3516872DB
    EVENT TIMESTAMP [55] 4 1365799643
    NAS PORT TYPE [61] 4 PPPoEoQinQ(34)
    NAS PORT ID [87] 11 3/2/2:1.100
    VSA [26] 29 DSL(3561)
      AGENT CIRCUIT ID [1] 13 circuit-id-12
      AGENT REMOTE ID [2] 12 remote-id-12
    VSA [26] 14 Alcatel(6527)
      SUBSC ID STR [11] 12 pppoe-user12
"
```

Next, the creation of the dynamic data service starts. As this is the first SAP for this service, the script which we reviewed above first creates the service instance.

```
12 2013/04/12 20:47:23.74 UTC MINOR: DEBUG #2001 Base dyn-script cli 1/1
"dyn-script cli 1/1: epipe:EPipe-CustomerName(cli 172 dict 0->31)

configure
  service
    epipe 1000 customer 1 create
      service-name "EPipe-CustomerName"
      description "EPipe-CustomerName"
      no shutdown
    exit
  exit
exit
"
```

Next, the SAP and the SDP are created within this service by the main function.

```
14 2013/04/12 20:47:23.74 UTC MINOR: DEBUG #2001 Base dyn-script cli 1/1
"dyn-script cli 1/1: business_epipe:3/2/2:1.1(cli 418 dict 0->308)

configure
  service
    epipe 1000
```

```
      sap 3/2/2:1.1 create
        description "CustomerName-Circuit-1"
        ingress
          qos 3
        exit
        egress
          qos 3
        exit
      exit
      spoke-sdp-fec 1000 fec 129 aii-type 2 create
        pw-template-bind 2
        saii-type2 64496:192.0.2.1:3333
        taii-type2 64496:192.0.2.5:3333
        no shutdown
      exit
    exit
  exit
exit
"
```

The service is created and is now active. As two RADIUS accounting destinations are configured in the dynamic services policy a RADIUS Accounting-Start message is sent to each destination to indicate the service is up.

```
16 2013/04/12 20:47:23.76 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Accounting-Request(4) 172.31.1.2:1813 id 252 len 294 vrid 1 pol radius-server-
policy-2
    STATUS TYPE [40] 4 Start(1)
    NAS IP ADDRESS [4] 4 192.0.2.1
    SESSION ID [44] 22 D6E559000000D4516872DB
    NAS PORT ID [87] 9 3/2/2:1.1
    DELAY TIME [41] 4 0
    NAS IDENTIFIER [32] 5 BNG-1
    EVENT TIMESTAMP [55] 4 1365799643
    MULTI SESSION ID [50] 22 D6E559000000D1516872DB
    USER NAME [1] 24 subscriber12@domain2.com
    VSA [26] 29 DSL(3561)
      AGENT CIRCUIT ID [1] 13 circuit-id-12
      AGENT REMOTE ID [2] 12 remote-id-12
    VSA [26] 117 Alcatel(6527)
      DYN SERV SCRIPT PARAMS [165] 115 business_epipe={'t':('EPipe-
CustomerName','CustomerName-Circuit-
1','3','3','64496','192.0.2.5','192.0.2.1','3333')}
"

15 2013/04/12 20:47:23.76 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Accounting-Request(4) 172.31.1.2:1813 id 251 len 294 vrid 1 pol radius-server-
policy-2
    STATUS TYPE [40] 4 Start(1)
    NAS IP ADDRESS [4] 4 192.0.2.1
    SESSION ID [44] 22 D6E559000000D4516872DB
    NAS PORT ID [87] 9 3/2/2:1.1
    DELAY TIME [41] 4 0
    NAS IDENTIFIER [32] 5 BNG-1
    EVENT TIMESTAMP [55] 4 1365799643
```

```
    MULTI SESSION ID [50] 22 D6E559000000D1516872DB
    USER NAME [1] 24 subscriber12@domain2.com
    VSA [26] 29 DSL(3561)
      AGENT CIRCUIT ID [1] 13 circuit-id-12
      AGENT REMOTE ID [2] 12 remote-id-12
    VSA [26] 117 Alcatel(6527)
      DYN SERV SCRIPT PARAMS [165] 115 business_epipe={'t':('EPipe-
CustomerName','CustomerName-Circuit-
1','3','3','64496','192.0.2.5','192.0.2.1','3333')}
"
```

For both RADIUS accounting destinations the interim accounting updates are also configured.

```
21 2013/04/12 20:51:46.69 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Accounting-Request(4) 172.31.1.2:1813 id 173 len 511 vrid 1 pol radius-server-
policy-1
    STATUS TYPE [40] 4 Interim-Update(3)
    NAS IP ADDRESS [4] 4 192.0.2.1
    SESSION ID [44] 22 D6E559000000D4516872DB
    NAS PORT ID [87] 9 3/2/2:1.1
    DELAY TIME [41] 4 0
    NAS IDENTIFIER [32] 5 BNG-1
    EVENT TIMESTAMP [55] 4 1365799906
    SESSION TIME [46] 4 125174
    MULTI SESSION ID [50] 22 D6E559000000D1516872DB
    USER NAME [1] 23 subscriber12@domain2.com
    VSA [26] 27 DSL(3561)
      AGENT CIRCUIT ID [1] 12 circuit-id-12
      AGENT REMOTE ID [2] 11 remote-id-12
    VSA [26] 241 Alcatel(6527)
      DYN SERV SCRIPT PARAMS [165] 115 business_epipe={'t':('EPipe-
CustomerName','CustomerName-Circuit-
1','3','3','64496','192.0.2.5','192.0.2.1','3333')}
      INPUT_INPROF_OCTETS_64 [19] 10 0x00010000000000000000
      INPUT_OUTPROF_OCTETS_64 [20] 10 0x00010000000000000000
      INPUT_INPROF_PACKETS_64 [23] 10 0x00010000000000000000
      INPUT_OUTPROF_PACKETS_64 [24] 10 0x00010000000000000000
      INPUT_HIGH_OCTETS_OFFER_64 [73] 10 0x00010000000000000000
      INPUT_LOW_PACK_OFFER_64 [76] 10 0x00010000000000000000
      INPUT_HIGH_PACK_OFFER_64 [75] 10 0x00010000000000000000
      INPUT_LOW_OCTETS_OFFER_64 [74] 10 0x00010000000000000000
      INPUT_UNC_PACK_OFFER_64 [78] 10 0x00010000000000000000
      INPUT_UNC_OCTETS_OFFER_64 [77] 10 0x00010000000000000000
      INPUT_HIGH_PACK_DROP_64 [71] 10 0x00010000000000000000
      INPUT_LOW_PACK_DROP_64 [72] 10 0x00010000000000000000
      INPUT_HIGH_OCTETS_DROP_64 [69] 10 0x00010000000000000000
      INPUT_LOW_OCTETS_DROP_64 [70] 10 0x00010000000000000000
      OUTPUT_INPROF_OCTETS_64 [21] 10 0x0001000000000000033c
    VSA [26] 84 Alcatel(6527)
      OUTPUT_OUTPROF_OCTETS_64 [22] 10 0x00010000000000000000
      OUTPUT_INPROF_PACKETS_64 [25] 10 0x0001000000000000000b
      OUTPUT_OUTPROF_PACKETS_64 [26] 10 0x00010000000000000000
      OUTPUT_INPROF_PACK_DROP_64 [81] 10 0x00010000000000000000
      OUTPUT_OUTPROF_PACK_DROP_64 [82] 10 0x00010000000000000000
      OUTPUT_INPROF_OCTS_DROP_64 [83] 10 0x00010000000000000000
      OUTPUT_OUTPROF_OCTS_DROP_64 [84] 10 0x00010000000000000000
```

```
"

19 2013/04/12 20:48:56.69 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Accounting-Request(4) 172.31.1.2:1813 id 253 len 241 vrid 1 pol radius-server-
policy-2
    STATUS TYPE [40] 4 Interim-Update(3)
    NAS IP ADDRESS [4] 4 192.0.2.1
    SESSION ID [44] 22 D6E559000000D4516872DB
    NAS PORT ID [87] 9 3/2/2:1.1
    DELAY TIME [41] 4 0
    NAS IDENTIFIER [32] 5 BNG-1
    EVENT TIMESTAMP [55] 4 1365799736
    SESSION TIME [46] 4 125004
    MULTI SESSION ID [50] 22 D6E559000000D1516872DB
    USER NAME [1] 23 subscriber12@domain2.com
    VSA [26] 27 DSL(3561)
      AGENT CIRCUIT ID [1] 12 circuit-id-12
      AGENT REMOTE ID [2] 11 remote-id-12
    VSA [26] 61 Alcatel(6527)
      DYN SERV SCRIPT PARAMS [165] 115 business_epipe={'t':('EPipe-
CustomerName','CustomerName-Circuit-
1','3','3','64496','192.0.2.5','192.0.2.1','3333')}
"
```

The “Stats-Type” in the dynamic service policy (or obtained via RADIUS in a VSA) defines what information is sent back to the accounting server (per server). In this example one was set to Stats-Type “time” and the other to “volume-time”. The first accounting message displays the content of “volume-time”. A full set of statistics counters per service class are provided for the dynamic service. This is equivalent to the extended accounting statistics also provided in the ESM context. The second accounting message shows the content of “time”. No volume statistics counters are provided in this case.

Once the dynamic data services are instantiated they can be displayed with the regular show commands.

```
A:BNG-1# show service service-using
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm  Opr  CustomerId Service Name
-------------------------------------------------------------------------------
1            VPLS      Up   Up   1          VPLS_For_Capture_SAPs
2            VPRN      Up   Up   1          VPRN_Control_Channel
3            VPRN      Up   Up   1          VPRN_REsidential_Subs
4            IES       Up   Up   1
10           VPRN      Up   Up   1
99           Mirror    Up   Up   1
500          Mirror    Up   Up   1
[1000]       Epipe     Up   Up   1          EPipe-CustomerName
[1001]       VPLS      Up   Up   1          VPLS-CustomerName
[1002]       IES       Up   Up   1          IES-CustomerName
[5000]       IES       Up   Up   1          IES-5000
```

```
[9999]        VPRN      Up   Up   1          VPRN-CustomerName
10001         VPLS      Up   Up   1
10002         Epipe     Up   Up   1
-snip-
-------------------------------------------------------------------------------
Matching Services : 20
-------------------------------------------------------------------------------
Dynamic Services : 5, indicated by [<svc-id>]
-------------------------------------------------------------------------------
===============================================================================
```

The dynamically created services are shown in the standard service list with their service IDs between brackets. It is possible to filter only the dynamic services using the **origin dyn-script** option.

```
A:BNG-1# show service service-using origin dyn-script
===============================================================================
Services
===============================================================================
ServiceId    Type      Adm  Opr  CustomerId Service Name
-------------------------------------------------------------------------------
[1000]        Epipe     Up   Up   1          EPipe-CustomerName
[1001]        VPLS      Up   Up   1          VPLS-CustomerName
[1002]        IES       Up   Up   1          IES-CustomerName
[5000]        IES       Up   Up   1          IES-5000
[9999]        VPRN      Up   Up   1          VPRN-CustomerName
-------------------------------------------------------------------------------
Matching Services : 5
-------------------------------------------------------------------------------
Dynamic Services : 5, indicated by [<svc-id>]
-------------------------------------------------------------------------------
===============================================================================
```

Similarly, the active SAPs can also be shown with the regular command.

```
A:BNG-1# show service sap-using
===============================================================================
Service Access Points
===============================================================================
PortId                          SvcId      Ing.   Ing.    Egr.   Egr.   Adm  Opr
                                           QoS    Fltr    QoS    Fltr
-------------------------------------------------------------------------------
3/2/1:*.100                     1          1      none    1      none   Up   Up
3/2/1:*.200                     1          1      none    1      none   Up   Up
3/2/2:*.100                     1          1      none    1      none   Up   Up
[3/2/1:4.100]                   2          1      none    1      none   Up   Up
[3/2/2:1.100]                   2          1      none    1      none   Up   Up
3/2/2:1000.1000                 2          1      none    1      none   Up   Up
[3/2/1:2.200]                   3          1      none    1      none   Up   Up
[3/2/1:3.200]                   3          1      none    1      none   Up   Up
3/2/1:1001.1001                 3          1      none    1      none   Up   Up
3/2/2:500.500                   3          1      none    1      none   Up   Up
3/2/2:100.100                   4          1      none    1      none   Up   Up
3/2/2:99.99                     99         1      none    1      none   Up   Up
[3/2/2:1.1]                     [1000]     3      none    3      none   Up   Up
```

```
[3/2/2:1.3]                     [1001]     3     none    3     none   Up   Up
[3/2/2:1.4]                     [1002]     5     ip4+ip6 6     ip4+i* Up   Up
[3/2/1:4.3]                     [5000]     1     none    1     none   Up   Up
[3/2/2:1.2]                     [9999]     3     ip4     3     ip4    Up   Up
3/2/1:99.99                     10001      1     none    1     none   Up   Up
3/2/19:100                      10001      1     none    1     none   Up   Up
3/2/20:100                      10001      1     none    1     none   Up   Up
-snip-
-------------------------------------------------------------------------------
Number of SAPs : 31
-------------------------------------------------------------------------------
Number of Managed SAPs : 4, indicated by [<sap-id>]
-------------------------------------------------------------------------------
Number of Dynamic Service SAPs : 5, indicated by [<sap-id>] [<svc-id>]
-------------------------------------------------------------------------------
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

The description at the end of this show command explains how the dynamic services SAPs are displayed. Note that there are managed SAPs created for the control channel as well as dynamic data services SAPs.

If only the SAPs for dynamic data services should be displayed, the command **show service sap-using dyn-script** can be used.

```
A:BNG-1# show service sap-using dyn-script
===============================================================================
Service Access Points
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
[3/2/2:1.1]                     [1000]     3     none    3     none   Up   Up
[3/2/2:1.3]                     [1001]     3     none    3     none   Up   Up
[3/2/2:1.4]                     [1002]     5     ip4+ip6 6     ip4+i* Up   Up
[3/2/1:4.3]                     [5000]     1     none    1     none   Up   Up
[3/2/2:1.2]                     [9999]     3     ip4     3     ip4    Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 5
-------------------------------------------------------------------------------
Number of Dynamic Service SAPs : 5, indicated by [<sap-id>] [<svc-id>]
-------------------------------------------------------------------------------
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

Dynamic data services CLI is not saved as part of the configuration file. The active dynamic data services configuration is hidden in the output of the "admin display-config" CLI command. To display the dynamic services configuration, use:

- "info" in a configuration CLI context for SR OS Releases prior to 14.0.R1
- "info include-dynamic" in a configuration context for SR OS Release 14.0.R1 or later

Conclusion

RADIUS-based dynamic data services provide an innovative way for business service provisioning. They are created both automatically and instantaneously.

Raw Formatting of DHCPv4/v6 Options in ESM

This chapter provides information about raw formatting of DHCPv4/v6 options in ESM.

Topics in this chapter include:

Applicability

This example is applicable to DHCPv4/v6 subscriber-hosts using the Routed Central Office ESM model on 7750 SR-7/12/12e, 7750 SR-c4/12 and 7450 ESS 6/7/12 in Mixed-Mode with IOM3-XP or IMM.

A local DHCPv4/v6 server is used for address/prefix assignment, which implies a 7x50 DHCP *relay* scenario (as opposed to a 7x50 DHCP *proxy* scenario where the IP address/prefix is assigned via a RADIUS server or an LUDB).

The configuration was tested in a single homed environment using SR OS 12.0.R4.

Overview

The 7x50 supports IP address assignment to its DHCP clients via two IP address assignment authorities:

- DHCP server — In this model the 7x50 behaves as a DHCP relay between the DHCP client and the 7x50 DHCP server.
- RADIUS/LUDB — In this model the IP address/prefix is assigned via a RADIUS server. or an LUDB and the 7x50 internal or external behaves as a proxy between the DHCP client and the non-DHCP aware RADIUS/LUDB.

Note that the term proxy can also refer to the functionality where the DHCP server is used for address assignment. In this case, the 7x50 would hide the DHCP server from the client and pretend to be the DHCP server to the client, passing the DHCP parameters between the client and the server (lease times, etc).

Within these two fundamental address assignment models, there are several mechanisms available on the 7x50 by which DHCP parameters (DHCP options and various parameters within the options) can be passed to the DHCP client during the address assignment phase.

For example, in the RADIUS/LUDB address assignment model, the DHCP parameters can be supplied via RADIUS, LUDB and Python, while in the DHCP server model, the DHCP parameters can also be also supplied via the DHCP server itself (in addition to RADIUS, LUDB and Python).

Some of the more commonly used DHCP parameters have their own RADIUS and CLI constructs. For example, a default router has its own RADIUS attribute(s):

```
Alc-Default-Router (26-6527-18)
```

or even its own CLI keyword:

```
config>router>dhcp>server>pool>subnet>options# default-router
config>service>router>dhcp>server>pool>subnet>options# default-router
config>subscr-mgmt>ludb>ipoe>host>options# default-router
```

Other less common DHCP options can be defined and inserted by the 7x50 DHCP relay agent using the pre-formatted (IP address, domain, or string) or the non-formatted (hex) custom-option CLI command:

```
config>router>dhcp>server>pool>options# custom-option
config>router>dhcp>server>pool>subnet>options# custom-option
config>router>dhcp6>server>pool>options# custom-option
config>router>dhcp6>server>pool>prefix>options# custom-option
config>service>vprn>dhcp>server>pool>options# custom-option
config>service>vprn>dhcp>server>pool>subnet>options# custom-option
config>service>vprn>dhcp6>server>pool>options# custom-option
config>service>vprn>dhcp6>server>pool>prefix>options# custom-option
config>subscriber-mgmt>ludb>ipoe>host>options# custom-option
```

The most flexible way of configuring DHCP parameters is by means of ‘raw’ (or hexadecimal) formatting. Any DHCP option can be hexadecimally (raw) formatted via the following RADIUS attributes:

```
Alc-ToClient-Dhcp-Options
Alc-ToClient-Dhcp6-Options
```

and/or via the custom-options CLI commands as outlined above. These options are then passed on to the DHCP client via the DHCP relay agent in the 7x50.

In addition to raw formatting via RADIUS or CLI, Python scripting can be used to intercept DHCP messages and modify their content.

The focus of this example is to demonstrate how the **raw** DHCP options are formatted via RADIUS. The messages can be optionally pre/post-processed by a Python script in the 7x50 before they are passed on to the DHCP client.

In this example, the following DHCP parameters are passed to the DHCP client using the **Alc-ToClient-Dhcp-Options** and the **Alc-ToClient-Dhcp6-Options** RADIUS attributes:

Table 44 **RADIUS Inserted Raw Options**

RADIUS	
DHCPv4 ToClient-Dhcp-Options	**DHCPv6 ToClient-Dhcp6-Options**
(default-)router [3] = 10.10.10.254	DNS server [23] = 2001:db8:1:1:1:1:1:1 2001:db8:1:1:1:1:1:2
DNS server [6] = 172.22.250.250 172.22.250/251	Domain search list [24] = ‘Nokia.com’ ‘test.com’
Domain name [15] = ‘alcatel.com’	Vendor specific-option [17] = ‘custom-test-option’
Custom option [230] = ‘custom test option’	
Renew time [58] = 5 min (300sec)	
Rebind time [59] = 6min 40sec (400sec)	

The DHCP parameters in the following DHCP messages are altered by a Python script:

Table 45 **Python Modified DHCP Fields**

Python	
DHCPv4 (DHCP Request)	**DHCPv6 (LDRA DHCP Request)**
Lease-time [51] = 8min 20sec (500sec)	IA-NA Preferred-Lifetime = 66min 40sec (4000sec)

Table 45 **Python Modified DHCP Fields (Continued)**

Python	
DHCPv4 (DHCP Request)	**DHCPv6 (LDRA DHCP Request)**
	IA-NA Valid-Lifetime = 66min 40sec (4000sec)
	IA-NA Renew-Time (T1) = 33min 20sec (2000sec)
	IA-NA Rebind-Time (T2) = 50min (3000sec)
	IA-PD Preferred-Lifetime = 66min 40sec (4000sec)
	IA-PD Valid-Lifetime = 66min 40sec (4000sec)
	IA-PD Renew-Time (T1) = 33min 20sec (2000sec)
	IA-PD Rebind-Time (T2) = 50min (3000sec)

The following DHCP parameters are configured via CLI in the 7x50 DHCPv4/v6 server:

Table 46 **CLI Inserted DHCP Options**

CLI DHCP Server Pool/prefix Options	
DHCPv4	**DHCPv6**
DNS server [6] = 172.22.250.253	DNS server [23] = 2001:db8:1:1:1:1:1:3
Custom option [231] = ‘dhcp injected custom option 231’	Custom option [232]= ‘v6 custom option 232’
	IA-NA Preferred-Lifetime = 20min (1200sec)
	IA-NA Valid-Lifetime = 20min (1200sec)
	IA-NA Renew-Time (T1) = 10 min (600 sec)
	IA-NA Rebind-Time (T2) = 15min (900 sec)
	IA-PD Preferred-Lifetime = 20min (1200sec)
	IA-PD Valid-Lifetime = 20min (1200sec)
	IA-PD Renew-Time (T1) = 10min (600 sec)

The RADIUS injected raw options are prepended by the DHCP relay agent in the 7x50 to any existing DHCP options already present in the DHCP message before being sent to the client. The existing options could be generated by the DHCP server (internal or external) or by the LUDB. No check is performed on the outgoing DHCP message towards the client in order to verify whether any of the RADIUS inserted

options are already present in the DHCP message. This could potentially lead to duplication of DHCP options in the outgoing DHCP messages in case that the same option is inserted via the DHCP server and via RADIUS. To prove the point, this example supplies the same DHCP option (with different values) via multiple sources (RADIUS and CLI).

Configuration of DHCP lease related times requires closer examination. In DHCPv4, the DHCP lease-time option (51) is always supplied by the DHCPv4 server (this cannot be disabled). In case the lease-time is also supplied via RADIUS in an Alc-ToClient-Dhcp-Options VSA, the client would receive two lease-times for the same IP address. This can lead to unpredictable behavior not only on the client side but also on the 7x50 DHCPv4 server side since the DHCPv4 server (and the 7x50 DHCPv4 relay agent) creates the lease state only for the lease-time supplied by the DHCP server, and ignores the one supplied via RADIUS or LUDB. This scenario is shown in Figure 210:

1. DHCP Discover arrives.
2. Radius authentication is triggered.
3. RADIUS returns lease-time value ‘B’ (Alc-ToClient-DHCP-Option) in Authentication-Accept message.
4. DHCP Discover is forwarded by the DHCP relay agent to the DHCP server.
5. DHCP server offers an IP lease with the configured lease-time of ‘A’.
6. The DHCP offer is sent to the DHCP relay agent.
7. The DHCP relay agent appends the lease-time ‘B’ supplied by RADIUS to the DHCP message.
8. The DHCP relay forwards the message to the DHCP client with both lease-times ‘A’ and ‘B’.

   Note that the example in Figure 210 does not represent a typical deployment case. This example is solely chosen to clarify the behavior in 7x50.

Figure 210 **DHCPv4 Lease-Time Inserted by RADIUS and DHCPv4 Server**

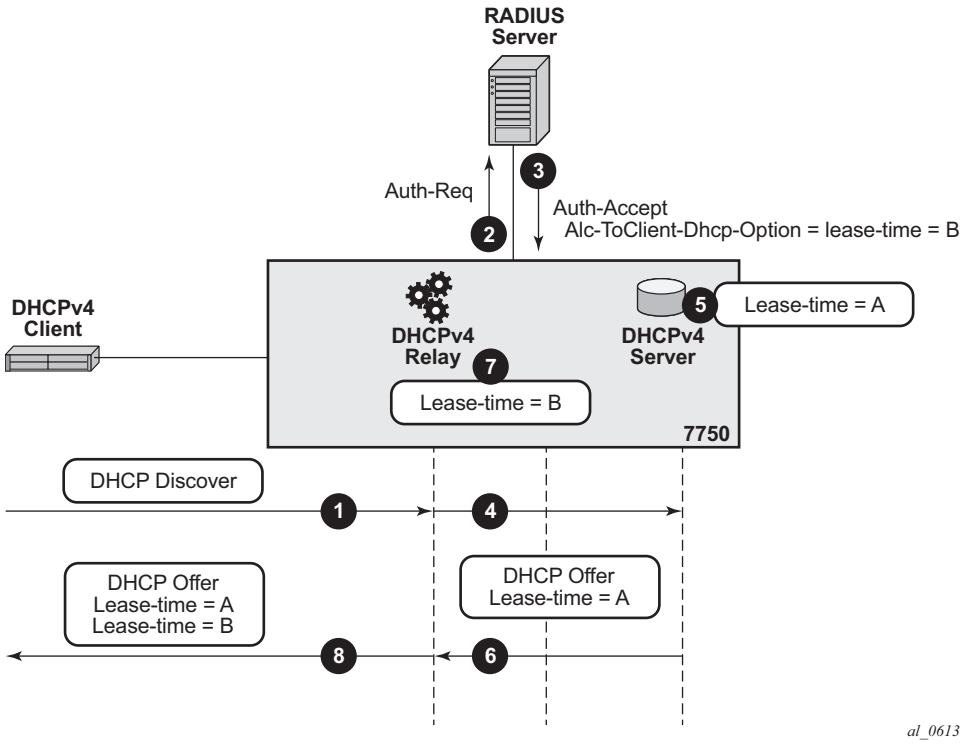

To ensure DHCPv4 lease time unambiguity, the lease-time should be supplied by a single source, in this case by the DHCPv4 server.

Since this eliminates RADIUS as a source of the DHCPv4 lease-time, an alternate method operating on the *raw level* is used to influence the automatic selection of the lease-time in the DHCPv4 server. This alternate method relies on the fact that the DHCPv4 server accepts hints received from the client as to what the desired lease-time should be. In other words, if the client sends the option 51 (lease-time) with a specific value, the 7x50 DHCPv4 server will honor this hint, as long as this value is within the configured range of values specified in the DHCP server. To demonstrate this behavior, a Python script is invoked upon receipt of a DHCPv4 Request message during the IP address assignment process (DORA – Discover-Offer-Request-Ack). The Python script inserts a new option 51 with the desired value for the lease-time. The DHCPv4 server honors this hint from the client and it returns the requested lease-time back to the client. This scenario is shown in Figure 211:.

1. DHCP Discover arrives.
2. DHCP Discover is intercepted by the Python processing engine and the lease-time ‘B’ is inserted in DHCP Discover message. This is then used as a hint to the DHCP server.

3. DHCP Discover message is sent to the DHCP relay agent.
4. RADIUS authentication is triggered.
5. User is authenticated. This time lease-time is not returned via RADIUS.
6. DHCP Discover is forwarded to the DHCP server.
7. The DHCP server honors the hint from the DHCP Discover and offers lease-time ‘B’, even though the server is configured with lease-time ‘A’.
8. The DHCP server replies with a DHCP Offer message.
9. DHCP Offer is forwarded by the DHCP relay agent to the client.

Figure 211 **Python Injected Hint for Lease-Time**

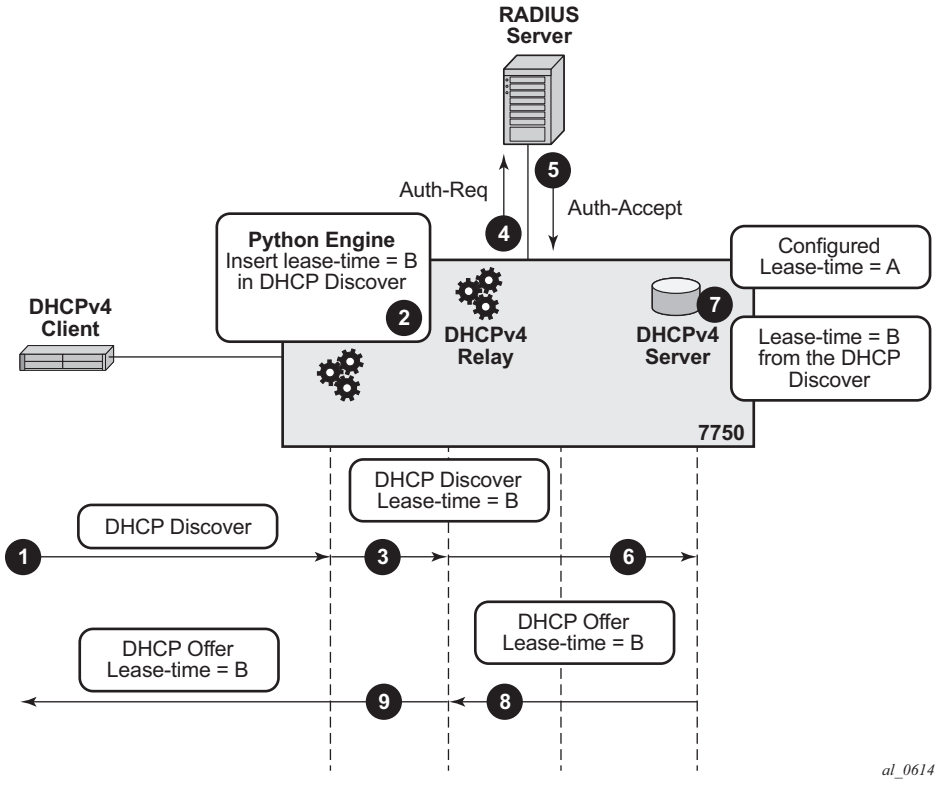

By default the local 7x50 DHCP server does **not** inject Renew (T1) and Rebind (T2) times so these two timers can still be supplied via RADIUS without duplication by the local 7x50 DHCP server.

When it comes to lease-time related parameters, the behavior of the 7x50 DHCPv6 server is different from the behavior of the 7x50 DHCPv4 server.

DHCPv6 lease related timers are **not** DHCP options. Instead, they are parameters within the IPv6 addressing option. An IPv6 address or prefix is assigned to the client via the IA-NA or IA-PD option, which contains additional parameters (which are not considered options) such as the IP address/prefix and the lease related timers. Figure 212 shows the IA-NA option that carries the T1/T2 parameters.

Figure 212 **Format of the IA-NA Option**

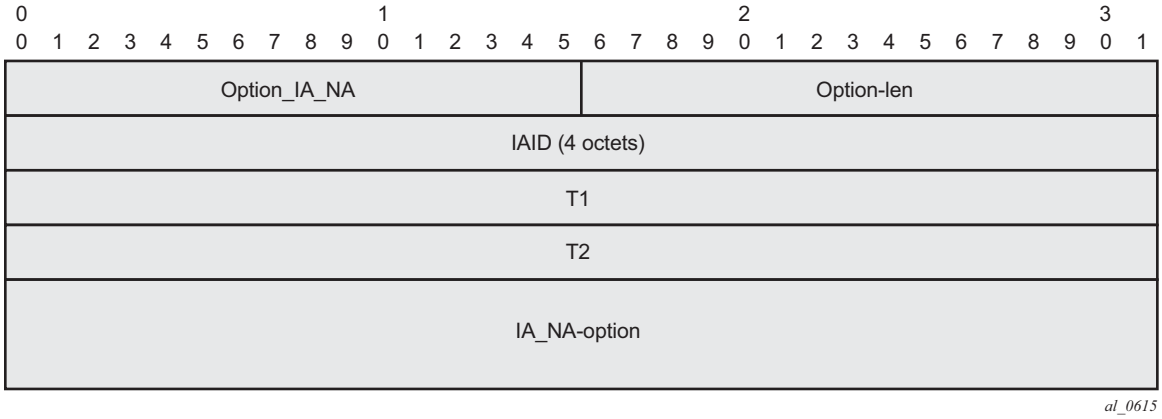

The format of the IA address option is shown in Figure 213. This option carries preferred and valid lifetimes.

Figure 213 **Format of the IA Address Option**

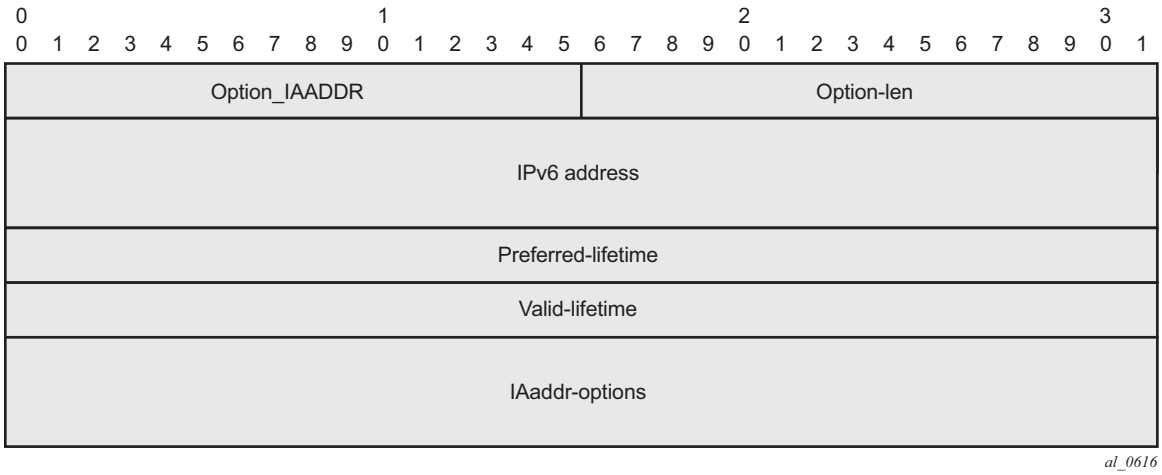

In this example, the IPv6 address/prefix is provided by the local 7x50 DHCPv6 server and as such, RADIUS cannot modify the parameters within the DHCPv6 options supplied by the DHCP server. Therefore, the desired IPv6 lease timers (preferred-life time, valid-lifetime, renew-time[T1], rebind-time[T2]) are part of the IPv6 pool configuration in the 7x50 DHCPv6 server.

Alternatively Python can be used to intercept the outgoing DHCPv6 message and then change the timers within the IA-NA and IA-PD options. Although this would configure the lease timers for the client, the action of modifying the outgoing DHCP6 message occurs after the DHCPv6 server processing. This would result in different lease times in the client and the DHCPv6 server, without any intermediary between them (such as a DHCPv6 Proxy) to deal with the differences.

For consistency purposes with the DHCPv4 example, a Python script processes the incoming DHCPv6 message (DHCPv6 Request) altering the lease timers (preferred/valid/renew/rebind) as a hint to the DHCPv6 server to request those values. However, the 7x50 DHCPv6 server does not honor those hints and uses its own values (default or configured) instead.

Configuration

The topology is shown in Figure 214.

Figure 214 **Topology**

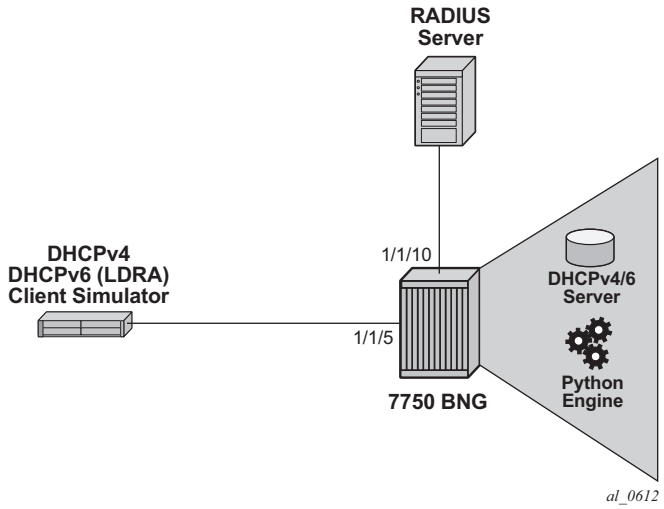

Access Ethernet Port with QinQ Encapsulation

```
configure port 1/1/5
    ethernet
            mode access
          encap-type qinq
    exit
     no shutdown
```

Capture SAP

A capture SAP is used to dynamically detect VLAN ID(s) in incoming DHCP (trigger) packets. This example uses RADIUS authentication along with Python scripting for DHCP message processing and therefore the authentication and Python policies must be configured under the capture SAP.

```
configure service vpls 10
    sap 1/1/5:1.* capture-sap create
        description "circuit-id authentication"
        trigger-packet dhcp dhcp6
        dhcp-python-policy "acg"
        dhcp6-python-policy "acg"
        authentication-policy "rad"
```

MSAP-Policy Configuration

The MSAP-policy defines the anti-spoofing mode which is set to next-hop MAC (nh-mac) in this example. It also defines the default subscriber management parameters in case that they are not supplied via LUDB or RADIUS.

MSAP-policy configuration is mandatory when a capture-SAP is deployed. In this example, the MSAP-policy name is supplied via RADIUS:

```
Alc-MSAP-Policy = "msaps"

configure subscriber-mgmt msap-policy “msaps”
    sub-sla-mgmt
        sub-ident-policy “sub_ident_pol”
        multi-sub-sap limit 500
    exit
    ies-vprn-only-sap-parameters
        anti-spoof nh-mac
    exit
```

Subscriber-Interface and Group-Interface Configuration

In this example the subscriber-interface is a ‘numbered interface’ in which the interface IPv4 address and the interface IPv6 prefixes are explicitly configured. The IPv4 address is used as the default-gateway by the IPoE attached clients. The IPv4 subnet to which this address belongs and the configured IPv6 prefixes are used for routing aggregation and are treated as local subnets/prefixes in the 7x50 routing table.

The managed (dynamic) SAPs are created under the group-interface which contains the reference to the authentication-policy name, the Python script, the v4/6 policy names and the DHCPv4/v6 relay related configuration settings (for example, a reference to DHCP servers). Both the authentication-policy name and the Python policy name referenced under the group-interface must match those configured under the capture-SAP.

```
configure service vprn 1
    subscriber-interface "int1-1" create
    address 10.10.10.254/24  # Numbered IPv4 subscriber interface.
    ipv6
        delegated-prefix-len 54
        subscriber-prefixes
            prefix 2001:db8:3::/48 pd       # Numbered IPv6 subscriber interface.
            prefix 2001:db8:4::/48 wan-host # Numbered IPv6 subscriber interface.
        exit
    exit
    group-interface "g1-1" create
        ipv6
            router-advertisements
                no shutdown
            exit
            dhcp6
                python-policy "acg"  # Python script for DHCPv6 messages.
                relay
                    server 2001:db8::1001 # IPv6 address of the DHCPv6 server.
                    client-applications dhcp
                    no shutdown
                exit
            exit
        exit
        dhcp
            python-policy "acg”               # Python script for DHCPv4 messages.
            option
                action keep            # Keep option82 in the received DHCP packet.
                vendor-specific-option
                    pool-name # Pool-
name obtained via RADIUS (or LUDB) will be passed
                              # via DHCP relay to the local DHCP server.  This name
                              # will be used for pool selection in DHCPv4 server.
                exit
            exit
            server 192.168.100.1 # IPv4 address of the DHCPv4 server.
            lease-populate 100   # Maximum number of DHCPv4 lease under each
                                 # SAP of the group-interface.
            client-applications dhcp
            no shutdown
        exit
        authentication-policy "rad" # RADIUS authentication policy.
    exit
exit
```

Loopback (DHCP) Interface Configuration

The loopback interface is used for the DHCPv4/v6 server binding. It is configured with the IPv4/IPv6 addresses which are referenced from the DHCP relay configuration under the group-interface.

```
configure service vprn 1
    interface "loopback1-1" create
        address 192.168.100.1/32        # IPv4 address of the DHCPv4 server.
        ipv6
            address 2001:db8::1001/128  # IPv6 address of the DHCPv6 server.
            local-dhcp-server "v6"      # Binding of the DHCPv6 server
                                        # to this interface.
        exit
        local-dhcp-server "v4"          # Binding of the DHCPv4 server
                                        # to this interface.
        loopback
    exit
```

DHCPv4/6 Server Configuration

The local DHCP server configuration contains the pool selection method, pool information and DHCP options which are passed to the DHCP client at IP address/ prefix assignment time.

```
configure service vprn 1
    dhcp
        local-dhcp-server "v4"
            use-pool-from-client  # Pool-name received in the DHCP messages
                                  # sent by the DHCP relay. The pool-name
                                  # is used in pool selection.
            pool "non-shared-left"
                options
                    dns-
server 172.22.250.253  # DHCPv4 option passed on to the client.
                    custom-
option 231 string "dhcp server injected custom option 231"
                exit                            # DHCPv4 option passed on to the clie
nt.
                subnet 10.10.10.0/24 create
                    address-
range 10.10.10.5 10.10.10.100 # IPv4 address range available
                                                                # for address allocation.
                exit
            exit
        exit
    exit
    dhcp6
        local-dhcp-server "v6"
            use-pool-from-client
            pool "pd-left" create
                options
```

```
                    dns-server 2001:db8:1:1:1:1:1:3
                    custom-option 232 string "v6 custom option 232"
                exit
                prefix 2001:db8:4::/
48 pd     # IPv6 prefix range available for delegated
                                            # prefix allocation by this DHCPv6 ser
ver.
                    preferred-lifetime min 20 # Preferred lifetime of the allocated
                                            # delegated prefix.
                    rebind-timer min 15       # Rebind (T2) time of the allocated
                                            # delegated prefix.
                    renew-timer min 10        # Renew (T1) time of the allocated
                                            # delegated prefix.
                    valid-lifetime min 20     # Valid lifetime of the allocated
                                            # delegated prefix.
                exit
            exit
            pool "wan-left" create
                options
                    dns-server 2001:db8:1:1:1:1:1:3
                    custom-option 232 string "v6 custom option 232"
                exit
                prefix 2001:db8:3::/56 wan-host
                    preferred-lifetime min 20 # Preferred lifetime of the
                                            # allocated IPv6 address.
                    rebind-timer min 15       # Rebind (T2) time of the
                                            # allocated IPv6 address.
                    renew-timer min 10        # Renew (T1) time of the
                                            # allocated IPv6 address.
                    valid-lifetime min 20     # Valid lifetime of the
                                            # allocated IPv6 address.
                exit
            exit
            no shutdown
        exit
    exit
```

RADIUS Authentication-Policy Configuration

The RADIUS authentication-policy is referenced under the capture-sap and under the group-interface configuration.

```
authentication-policy "rad" create
    password "ALU" hash2
     radius-authentication-server
     router "Base"
         server 1 address 192.168.114.1 secret "ALU" hash2
     exit
     user-name-format circuit-id
     include-radius-attribute
        circuit-id
        remote-id
        nas-port-id
        nas-identifier
    exit
exit
```

Subscriber-Identification Policy

The subscriber-identification policy in this example defines a mapping method between the subscriber strings and the predefined subscriber profiles (*sub* and *sla*) locally configured on the 7x50. In our example the subscriber strings (*sub* and *sla*) are provided via RADIUS and are directly mapped to the preconfigured sub-profiles and sla-profiles with the matching names.

The subscriber-identification policy can be configured with default subscriber profiles in case the strings are not explicitly obtained via other means (RADIUS, LUSB, Python or statically provisioned). Subscriber-identification policy configuration is mandatory.

```
sub-ident-policy "sub_ident_pol" create
     sub-profile-map
         use-direct-map-as-default
     exit
     sla-profile-map
         use-direct-map-as-default
     exit
```

Sla-Profile and Sub-Profile Configuration

The following is the configuration of the sub-profile and the sla-profile which are used to setup the subscriber-host. The sla and sub profiles are mandatory when creating subscriber-hosts in 7x50.

```
sla-profile "sla-profile-1" create
    ingress
        qos 2
         exit
    exit
     egress
         qos 2
         exit
     exit
exit

sub-profile "sub-profile-1" create
exit
```

Python-Policy Configuration

The python-policy defined below is applied under the capture-sap and under the group-interface. It references the python-script command which defines the location of the script. A python-policy specifies the DHCP messages along with the direction to which the script processing applies.

The DHCPv4 script in this example is applied to incoming DHCPv4 Request messages. The python script inserts the lease-time option in the DHCPv4 Request message as a hint to the DHCPv4 server.

Similar logic is applied to the incoming Lightweight DHCPv6 Relay Agent

(LDRA) DHCPv6 messages where IA-NA and IA-PD related lease times are altered. Note that in the DHCPv6 case the local DHCPv6 server does not honor the hint and therefore the lease related times are explicitly configured in the DHCPv6 server.

```
python-script "acg" create
     action-on-
fail passthrough #In case of python script failure, do not drop the
                 message but instead continue with message processing
                   in 7750.
     primary-url "ftp://a.b.c.d/pub/configs/alu/SIMS/acg.py"
     no shutdown
exit
python-script "acg6" create
    action-on-fail passthrough
    primary-url "ftp://a.b.c.d/pub/configs/alu/SIMS/acg6.py"
    no shutdown
exit
python-policy "acg" create  #Python policy that is applied under the capture-sap and
                             under the group-interface.
    dhcp request direction ingress script "acg"
    dhcp6 relay-forward direction ingress script "acg6"
exit
```

Python Script Configuration

In this example the Python script is located in an external location and downloaded to the 7x50 once the python-script CLI node is enabled (**no shutdown**).

The DHCPv4 Python script has exception code included (try/except statements). This makes script debugging easier in case one of the commands in the script fails.

For simplicity reasons, the exception code is removed from the DHCPv6 Python script. Note that in real deployments it is recommended for the exception code to be included in all Python scripts.

DHCPv4 Python Script:

```
from alc import dhcpv4
try:
    myopt = dhcpv4.getOptionList()
    if myopt != []:
        print "option-list ", repr(myopt)
         print "\n"
except Exception:
    print "Can't retrieve DHCP options"
#lease 500s 8min 20sec
try:
    dhcpv4.set(51,('\x00\x00\x01\xf4', #Insert the lease-
time (opt51) in the incoming
                                        DHCPv4 request as a hint to the DHCPv4 server
.
except Exception:
    print "Can't set time lease"
```

DHCPv6 Python Script:

```
from alc import dhcpv6
import struct
packet = dhcpv6.get_relaymsg()# Extract the original DHCPv6 packet within LDRA.

msgType = ord(packet.msg_type) # Get the message type.
ia_na = packet.get_iana()    # Store the IA-
NA option for further processing later on.
ia_pd = packet.get_iapd()    # Store the IA-
PD option for further processing later on.



if msgType == 3: # If the message in the LDRA packet is DHCPv6 Request, insert the l
ease related times in address/prefic options.

    ia_na[0][1] = '\x00\x00\x07\xd0'# Set the renew time (T1) in IA-NA to 2000sec.
    ia_na[0][2] = '\x00\x00\x0b\xb8'# Set the rebind time (T2) in IA-NA to 3000sec.
    ia_na[0][3][5][0][1] = '\x00\x00\x0f\xa0’  # Set the preferred time in IA-NA to
                                               # 4000sec.
    ia_na[0][3][5][0][2] = '\x00\x00\x0f\xa0'# Set the valid time in IA-
NA to 4000sec.
    packet.set_iana(ia_na) # Update the stored packet with the new values for IA-NA.

    ia_pd[0][1] = '\x00\x00\x07\xd0'# Set the renew time (T1) in IA-PD to 2000sec.
    ia_pd[0][2] = '\x00\x00\x0b\xb8'# Set the rebind time (T2) in IA-PD to 3000sec.
    ia_pd[0][3][26][0][0] = '\x00\x00\x0f\xa0’ # Set the preferred time in IA-PD to
                                               # 4000sec.
    ia_pd[0][3][26][0][1] = '\x00\x00\x0f\xa0'# Set the valid time in IA-
PD to 4000sec.
    packet.set_iapd(ia_pd) # Update the stored packet with the new values for IA-PD.
    dhcpv6.set_relaymsg(packet) # Insert the packet in the LDRA message.
```

RADIUS Access-Accept

Upon authentication, RADIUS returns the Access-Accept message with the following attributes:

```
Sending Access-Accept of id 66 to 192.168.114.2 port 64384
    Alc-Subsc-Prof-Str = "sub-profile-1"
    Alc-SLA-Prof-Str = "sla-profile-2"
    Alc-MSAP-Interface = "g1-1"
    Alc-MSAP-Policy = "msaps"
    Alc-MSAP-Serv-Id = 1
    Framed-Pool = "non-shared-left"
    Framed-IPv6-Pool = "wan-left"
    Alc-Delegated-IPv6-Pool = "pd-left"
    Alc-ToClient-Dhcp-Options += 0x03040a0a0afe
    Alc-ToClient-Dhcp-Options += 0x0608ac16fafaac16fafb
    Alc-ToClient-Dhcp-Options += 0x0f0b616c636174656c2e636f6d
    Alc-ToClient-Dhcp-Options += 0xe612637573746f6d2074657374206f7074696f6e
    Alc-ToClient-Dhcp-Options += 0x3a040000012c
    Alc-ToClient-Dhcp-Options += 0x3b0400000190
    Alc-ToClient-Dhcp6-
Options +=         0x0011001a0000197f00e60012637573746f6d2074657374206f7074696f6e
    Alc-ToClient-Dhcp6-
Options +=         0x0017002020010db80001000100010001000100120010db8000100010001000
100010002
    Alc-ToClient-Dhcp6-
Options +=         0x0018001e0e616c636174656c2d6c7563656e7403636f6d00047465737403636
f6d
```

It is possible to concatenate multiple DHCP options in a single RADIUS Alc-ToClient-DHCP6-Option but for clarity each option is in a separate attribute in this example.

The following table contains the explanation of the DHCP options inserted via RADIUS:

Table 47 **DHCP options inserted via RADIUS**

Alc-ToClient-Dhcp-Options += 0x03040a0a0afe (default) router (3) = 10.10.10.254
Alc-ToClient-Dhcp-Options += 0x0608ac16fafaac16fafb dns (6) = 172.16.250.250 172.16.250.251
Alc-ToClient-Dhcp-Options += 0x0f0b616c636174656c2e636f6d domain-name (15) = alcatel.com
Alc-ToClient-Dhcp-Options += 0xe612637573746f6d2074657374206f7074696f6e custom -option (230) = "custom test option"
Alc-ToClient-Dhcp-Options += 0x3a040000012c renewal time T1 (58) = 300s (5min)

Table 47 **DHCP options inserted via RADIUS (Continued)**

Alc-ToClient-Dhcp-Options += 0x3b0400000190 rebind time T2 (59) = 400s (6min 40sec)
Alc-ToClient-Dhcp6-Options += 0x0011001a0000197f00e60012637573746f6d2074657374206f7074696f6e v6 vendor option (17) [opt-id(2) len(2) entp-id(4) vopt-code(2) vlen(2) vdata] = 17 26 6527 230 18 "custom test option"
Alc-ToClient-Dhcp6-Options += 0x0017002020010db80001000100010001000100012001 0db80001000100010001000100010002 dns servers (23) [opt-id(2) len(2) servers-v6@] = 23 32 2001:0db8:0001:0001:0001:0001:0001:0001 2001:0db8:0001:0001:0001:0001:0001:0002
Alc-ToClient-Dhcp6-Options += 0x0018001e0e616c636174656c2d6c7563656e7403636f6d00047465737403636f6d0 domain list (24) = Nokia.com test.com [formatting as described in section 3.1 of RFC 1035 (as referenced by RFC 4704 and RFC 3315)].

Results and Verification

The results are verified via debug output and show commands on the 7x50, and also via pcap (Wireshark® packet capture) files on the DHCP client side.

Debug output on the 7x50 is enabled for DHCPv4/6 messages and for the Python script.

The DHCP debug output shows the options sent to the client in the DHCPv4/6 Ack/ Reply messages.

The following commands enables debugging information to be sent to the current telnet/ssh session:

```
*A:BNG1# configure  log
*A:BNG1>config>log# info
----------------------------------------------
        log-id 50 # Capturing and displaying debug output is configured via log.
            from debug-trace # Capture debug output.
            to session     # Output the debug to the current tcp/ssh session.
        exit
----------------------------------------------
```

The following commands enable DHCP related debugging:

```
*A:BNG1>config>log# show debug
debug
    router "1"
```

```
        ip
            dhcp
                detail-level high
                mode egr-ingr-and-dropped
            exit
            dhcp6
                mode egr-ingr-and-dropped
                detail-level high
            exit
        exit
        local-dhcp-server "v4"
            detail-level high
            mode egr-ingr-and-dropped
        exit
        local-dhcp-server "v6"
            detail-level high
            mode egr-ingr-and-dropped
        exit
    exit
```

DHCPv4 Results

The following output displays the DHCPv4 Request message as it was received by the 7x50 DHCP server.

This message has been modified by the Python script on ingress and the lease-time option [51] has been inserted as a hint to the DHCPv4 server.

Option [82] is partially added by the **access-node** (relay-agent —> circuit-id and remote-id) and partially by the internal 7x50 DHCP-relay (pool name).

```
32830 2014/07/24 03:02:46.44 WEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  v4
Rx DHCP Request

   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 10.10.10.254
   chaddr: 00:00:65:01:03:01    xid: 0x159dd536

   DHCP options:
   [82] Relay agent information: len = 42
      [1] Circuit-id: ds-left
      [2] Remote-id: remote0
      [9] Vendor-Specific info: len = 22
          Enterprise [6527] : len = 17
          [13] dhcpPool: non-shared-left
   [53] Message type: Request
   [54] DHCP server addr: 192.168.100.1
   [50] Requested IP addr: 10.10.10.34
   [51] Lease time: 500
   [255] End
```

The next output captures the DHCPv4 ACK message (within the 7x50) that is on its way to the client.

It can be observed that the DHCPv4 server inserted options are listed first:

- Opt[82] is echoed back by 7x50 DHCPv4 server
- Opt[53], [54], [51] and [1] are by default inserted by the local 7x50 DHCPv4 server and they cannot be disabled. The value for the lease-time [51] is set by the Python script.
- The next two options ([6] and [231]) are the options configured explicitly in the DHCPv4 server (Table 46).

The remaining options (with the exception of the *end* [255] option) are provided by RADIUS and they appear in the exact same order as they appear in the RADIUS Alc-ToClient-Dhcp-Options attributes (Table 44).

There are two options [6] since they are inserted by both DHCP and RADIUS server.

Custom options [231] and [230] are decoded in Table 44 and Table 46.

```
32834 2014/07/24 03:02:46.44 WEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 11 (g1-1),
   transmitted DHCP Boot Reply to Interface g1-1 (1/1/5:1.3) Port 68

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.10.10.34
   siaddr: 192.168.100.1     giaddr: 10.10.10.254
   chaddr: 00:00:65:01:03:01    xid: 0x159dd536

   DHCP options:
   [82] Relay agent information: len = 18
      [1] Circuit-id: ds-left
      [2] Remote-id: remote0
   [53] Message type: Ack
   [54] DHCP server addr: 192.168.100.1
   [51] Lease time: 500
   [1] Subnet mask: 255.255.255.0
   [6] Domain name server: 172.22.250.253
   [231] Unknown option: len = 38, value = 64 68 63 70 20 73 65 72 76 65 72
   20 69 6e 6a 65 63 74 65 64 20 63 75 73 74 6f 6d 20 6f 70 74 69 6f 6e 20 32
   33 31
   [3] Router: 10.10.10.254
   [6] Domain name server: length = 8
             172.22.250.250
             172.22.250.251
   [15] Domain name: alcatel.com
   [230] Unknown option: len = 18, value = 63 75 73 74 6f 6d 20 74 65 73 74
   20 6f 70 74 69 6f 6e
   [58] Renew timeout: 300
   [59] Rebind timeout: 400
   [255] End
```

The Wireshark® output shown on the next page is captured at the client side (N2X Ixia) and it effectively mirrors what is shown in the 7x50 debug output.

```
⊟ Bootstrap Protocol
    Message type: Boot Reply (2)
    Hardware type: Ethernet (0x01)
    Hardware address length: 6
    Hops: 0
    Transaction ID: 0x159dd536
    Seconds elapsed: 0
  ⊞ Bootp flags: 0x8000 (Broadcast)
    Client IP address: 0.0.0.0 (0.0.0.0)
    Your (client) IP address: 10.10.10.34 (10.10.10.34)
    Next server IP address: 192.168.100.1 (192.168.100.1)
    Relay agent IP address: 10.10.10.254 (10.10.10.254)
    Client MAC address: NetworkG_01:03:01 (00:00:65:01:03:01)
    Client hardware address padding: 00000000000000000000
    Server host name not given
    Boot file name not given
    Magic cookie: DHCP
  ⊟ Option: (53) DHCP Message Type
      Length: 1
      DHCP: ACK (5)
  ⊞ Option: (54) DHCP Server Identifier
  ⊟ Option: (51) IP Address Lease Time
      Length: 4
      IP Address Lease Time: (500s) 8 minutes, 20 seconds
  ⊞ Option: (1) Subnet Mask
  ⊟ Option: (6) Domain Name Server
      Length: 4
      Domain Name Server: 172.22.250.253 (172.22.250.253)
  ⊟ Option: (231) Private
      Length: 38
      Value: 6468637020736572766572206 6e6a656374656420637 3...
  ⊞ Option: (82) Agent Information Option
  ⊟ Option: (3) Router
      Length: 4
      Router: 10.10.10.254 (10.10.10.254)
  ⊟ Option: (6) Domain Name Server
      Length: 8
      Domain Name Server: 172.22.250.250 (172.22.250.250)
      Domain Name Server: 172.22.250.251 (172.22.250.251)
  ⊟ Option: (15) Domain Name
      Length: 11
      Domain Name: alcatel.com
  ⊟ Option: (230) Private
      Length: 18
      Value: 637573746f6d2074657374206f7074696f6e
  ⊟ Option: (58) Renewal Time Value
      Length: 4
      Renewal Time Value: (300s) 5 minutes
  ⊟ Option: (59) Rebinding Time Value
      Length: 4
      Rebinding Time Value: (400s) 6 minutes, 40 seconds
  ⊞ Option: (255) End
```

The show command for the DHCP-relay lease state only displays the well known options inserted by the DHCPv4 server. The custom option inserted by the DHCPv4 server and any of the RADIUS supplied options are not kept as part of the DHCP-relay lease state.

```
*A:BNG1# show service  id 1 dhcp lease-state detail
===============================================================================
DHCP lease states for service 1
===============================================================================
Service ID           : 1
IP Address           : 10.10.10.34
Client HW Address    : 00:00:65:01:03:01
Subscriber-interface : int1-1
Group-interface      : g1-1
SAP                  : [1/1/5:1.3]
Up Time              : 0d 00:10:46
Remaining Lease Time : 0d 00:07:35
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "ds-left"
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-2"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : ""
Category-Map-Name    : ""

Lease Info origin    : DHCP

Ip-Netmask           : 255.255.255.0
Broadcast-Ip-Addr    : N/A
Default-Router       : N/A
Primary-Dns          : 172.22.250.253
Secondary-Dns        : N/A
Primary-Nbns         : N/A
Secondary-Nbns       : N/A

ServerLeaseStart     : 07/24/2014 03:02:46
ServerLastRenew      : 07/24/2014 03:12:46
ServerLeaseEnd       : 07/24/2014 03:21:06
Session-Timeout      : N/A
Lease-Time           : 0d 00:08:20
DHCP Server Addr     : 192.168.100.1

Relay Agent Information
  Circuit Id         : ds-left
  Remote Id          : remote0
Radius User-Name     : "ds-left"
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
```

DHCPv6 Results

The DHCPv6 server receives the DHCPv6 Request message with Python modified lease times (preferred, valid, renew and rebind) for IA-NA and IA-PD.

```
32877 2014/07/24 03:15:28.32 WEST MINOR: DEBUG #2001 vprn1 DHCP server
"DHCP server:  v6
Rx DHCPv6 RELAY_FORW
     Hop Count : 1
     Link Addr : 2001:db8:4::
     Peer Addr : fe80::200:65ff:fe01:301
     Option : RELAY_MSG (9), Length : 184
       Msg Type : RELAY_FORW (12)
       Hop Count : 0
       Link Addr : ::
       Peer Addr : fe80::200:65ff:fe01:301
       Option : INTERFACE_ID (18), Length : 7
         Interface Id : 64732d6c656674 (ds-left)
       Option : RELAY_MSG (9), Length : 135
         Msg Type : REQUEST (3)
         Trans Id : 0x060000
         Option : ELAPSED_TIME (8), Length : 2
           Time : 0 seconds
         Option : CLIENTID (1), Length : 10
           LL  : HwTyp=0001,LL=000065010301
           0003000100065010301
         Option : SERVERID (2), Length : 10
           LL  : HwTyp=0001,LL=d896ff000000
           00030001d896ff000000
         Option : ORO (6), Length : 4
           Requested Option : IA_NA (3)
           Requested Option : IA_PD (25)
         Option : IA_NA (3), Length : 40
           IAID : 0
           Time1: 2000 seconds
           Time2: 3000 seconds
           Option : IAADDR (5), Length : 24
             Address : 2001:db8:3:1::1
             Preferred Lifetime : 4000 seconds
             Valid Lifetime     : 4000 seconds
         Option : IA_PD (25), Length : 41
           IAID : 0
           Time1: 2000 seconds
           Time2: 3000 seconds
           Option : IAPREFIX (26), Length : 25
             Prefix : 2001:db8:4:400::/54
             Preferred Lifetime : 4000 seconds
             Valid Lifetime     : 4000 seconds
     Option : VENDOR_OPTS (17), Length : 37
       Enterprise : 0000197f
       Option : WAN_POOL (1), Length : 8
         wan-left
       Option : PFX_POOL (2), Length : 7
         pd-left
       Option : PFX_LEN (3), Length : 1
```

The **hinted** DHCPv6 lease-times are not honored by the 7x50 DHCPv6 server and instead the 7x50 DHCPv6 server default values are inserted in the outgoing DHCPv6 Reply message towards the client as shown in the output below.

The explicitly configured DHCPv6 options are inserted by the DHCPv6 server first (Table 46) followed by the RADIUS supplied options inserted by the DHCPv6 relay (Table 44).

There are two DNS options [23] since they are supplied via two sources (DHCPv6 server and RADIUS Alc-ToClient-DHCP-Option VSA).

```
32885 2014/07/24 03:15:28.32 WEST MINOR: DEBUG #2001 vprn1 TIP
"TIP: DHCP6_PKT
   Outgoing DHCP6 Msg : RELAY_REPLY (13)
   to itf g1-1
     Hop Count : 0
     Link Addr : ::
     Peer Addr : fe80::200:65ff:fe01:301
     Option : RELAY_MSG (9), Length : 265
       Msg Type : REPLY (7)
       Trans Id : 0x060000
       Option : SERVERID (2), Length : 10
         LL  : HwTyp=0001,LL=d896ff000000
         00030001d896ff000000
       Option : CLIENTID (1), Length : 10
         LL  : HwTyp=0001,LL=000065010301
         0003000100065010301
       Option : IA_NA (3), Length : 40
         IAID : 0
         Time1: 600 seconds
         Time2: 900 seconds
         Option : IAADDR (5), Length : 24
           Address : 2001:db8:3:1::1
           Preferred Lifetime : 1200 seconds
           Valid Lifetime     : 1200 seconds
       Option : IA_PD (25), Length : 41
         IAID : 0
         Time1: 600 seconds
         Time2: 900 seconds
         Option : IAPREFIX (26), Length : 25
           Prefix : 2001:db8:4:400::/54
           Preferred Lifetime : 1200 seconds
           Valid Lifetime     : 1200 seconds
       Option : DNS_NAME_SRVR (23), Length : 16
         Server : 2001:db8:1:1:1:1:1:3
       Option : UNKNOWN (232), Length : 20
         7636206375737 46f6d206f7074696f6e20323332
       Option : VENDOR_OPTS (17), Length : 26
         Enterprise : 0000197f
         Option : UNKNOWN (230), Length : 18
           637573746f6d2074657374206f7074696f6e
       Option : DNS_NAME_SRVR (23), Length : 32
         Server : 2001:db8:1:1:1:1:1:1
         Server : 2001:db8:1:1:1:1:1:2
       Option : DOM_SRCH_LIST (24), Length : 30
         SearchList : .Nokia.com..test.com.
```

```
Option : INTERFACE_ID (18), Length : 7
  Interface Id : 64732d6c656674 (ds-left)
```

The Wireshark® capture of the DHCPv6 Reply message on the client side mirrors the debug information captured by the 7x50:

```
DHCPv6
  Message type: Relay-reply (13)
  Hopcount: 0
  Link address: :: (::)
  Peer address: fe80::200:65ff:fe01:301 (fe80::200:65ff:fe01:301)
  Relay Message
    Option: Relay Message (9)
    Length: 265
    Value: 070600000002000a00030001d896ff0000000001000a0003...
    DHCPv6
      Message type: Reply (7)
      Transaction ID: 0x060000
      Server Identifier
        Option: Server Identifier (2)
        Length: 10
        Value: 00030001d896ff000000
        DUID: 00030001d896ff000000
        DUID Type: link-layer address (3)
        Hardware type: Ethernet (1)
        Link-layer address: d8:96:ff:00:00:00
      Client Identifier
        Option: Client Identifier (1)
        Length: 10
        Value: 00030001000065010301
        DUID: 00030001000065010301
        DUID Type: link-layer address (3)
        Hardware type: Ethernet (1)
        Link-layer address: 00:00:65:01:03:01
      Identity Association for Non-temporary Address
        Option: Identity Association for Non-temporary Address (3)
        Length: 40
        Value: 0000000000000258000003840005001820010db800030001...
        IAID: 00000000
        T1: 600
        T2: 900
        IA Address
          Option: IA Address (5)
          Length: 24
          Value: 20010db80003000100000000000000001000004b0000004b0
          IPv6 address: 2001:db8:3:1::1 (2001:db8:3:1::1)
          Preferred lifetime: 1200
          Valid lifetime: 1200
```

```
⊟ Identity Association for Prefix Delegation
    Option: Identity Association for Prefix Delegation (25)
    Length: 41
    Value: 0000000000000258000003840001a0019000004b0000004b0...
    IAID: 00000000
    T1: 600
    T2: 900
  ⊟ IA Prefix
      Option: IA Prefix (26)
      Length: 25
      Value: 000004b0000004b03620010db80004040000000000000000...
      Preferred lifetime: 1200
      Valid lifetime: 1200
      Prefix length: 54
      Prefix address: 2001:db8:4:400:: (2001:db8:4:400::)
⊟ DNS recursive name server
    Option: DNS recursive name server (23)
    Length: 16
    Value: 20010db8000100010001000100010003
    DNS server address: 2001:db8:1:1:1:1:1:3 (2001:db8:1:1:1:1:1:3)
⊟ DHCP option 232
    Option: Unknown (232)
    Length: 20
    Value: 763620637573746f6d206f7074696f6e20323332
⊟ Vendor-specific Information
    Option: Vendor-specific Information (17)
    Length: 26
    Value: 0000197f00e60012637573746f6d2074657374206f707469...
    Enterprise ID: Panthera Networks, Inc. (6527)
  ⊞ option
⊟ DNS recursive name server
    Option: DNS recursive name server (23)
    Length: 32
    Value: 20010db8000100010001000100010001200010db800010001...
    DNS server address: 2001:db8:1:1:1:1:1:1 (2001:db8:1:1:1:1:1:1)
    DNS server address: 2001:db8:1:1:1:1:1:2 (2001:db8:1:1:1:1:1:2)
⊟ Domain Search List
    Option: Domain Search List (24)
    Length: 30
    Value: 0e616c636174656c2d6c7563656e7403636f6d0004746573...
    DNS Domain Search List
    Domain: alcatel-lucent.com
    Domain: test.com
⊟ Interface-Id
    Option: Interface-Id (18)
    Length: 7
    Value: 64732d6c656674
    Interface-ID: ds-left
```

The following command captures the information kept in the 7x50 DHCPv6 relay lease state:

```
*A:BNG1# show service  id 1 dhcp6 lease-state detail
===============================================================================
DHCP lease states for service 1
===============================================================================
Service ID           : 1
IP Address           : 2001:db8:3:1::1/128
Client HW Address    : 00:00:65:01:03:01
Subscriber-interface : int1-1
```

```
Group-interface      : g1-1
SAP                  : [1/1/5:1.3]
Up Time              : 0d 00:02:41
Remaining Lease Time : 0d 00:17:18
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "ds-left"
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-2"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : ""
Category-Map-Name    : ""
Dhcp6 ClientId (DUID): 00030001000065010301
Dhcp6 IAID           : 0
Dhcp6 IAID Type      : non-temporary
Dhcp6 Client Ip      : fe80::200:65ff:fe01:301
Primary-Dns          : N/A
Secondary-Dns        : N/A
Pool Name            : "wan-left"
Dhcp6 Server Addr    : 2001:db8::1001
Dhcp6 ServerId (DUID): 00030001d896ff000000
Dhcp6 InterfaceId    : ds-left
Dhcp6 RemoteId       : N/A

Lease Info origin    : DHCP

ServerLeaseStart     : 07/24/2014 03:15:28
ServerLastRenew      : 07/24/2014 03:15:28
ServerLeaseEnd       : 07/24/2014 03:35:27
Session-Timeout      : N/A
Radius User-Name     : "ds-left"
-------------------------------------------------------------------------------
Service ID           : 1
IP Address           : 2001:db8:4:400::/54
Client HW Address    : 00:00:65:01:03:01
Subscriber-interface : int1-1
Group-interface      : g1-1
SAP                  : [1/1/5:1.3]
Up Time              : 0d 00:02:41
Remaining Lease Time : 0d 00:17:18
Remaining SessionTime: N/A
Persistence Key      : N/A

Sub-Ident            : "ds-left"
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-2"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : ""
Category-Map-Name    : ""
Dhcp6 ClientId (DUID): 00030001000065010301
Dhcp6 IAID           : 0
Dhcp6 IAID Type      : prefix
Dhcp6 Client Ip      : fe80::200:65ff:fe01:301
Primary-Dns          : N/A
Secondary-Dns        : N/A
Pool Name            : "pd-left"
```

```
Dhcp6 Server Addr    : 2001:db8::1001
Dhcp6 ServerId (DUID): 00030001d896ff000000
Dhcp6 InterfaceId    : ds-left
Dhcp6 RemoteId       : N/A

Lease Info origin    : DHCP

ServerLeaseStart     : 07/24/2014 03:15:28
ServerLastRenew      : 07/24/2014 03:15:28
ServerLeaseEnd       : 07/24/2014 03:35:27
Session-Timeout      : N/A
Radius User-Name     : "ds-left"
-------------------------------------------------------------------------------
Number of lease states : 2
===============================================================================
```

Python Debug Output

DHCPv4

For debugging purpose a line is added to the Python script printing all DHCP option numbers present in the incoming DHCP packets.

It can also be observed that all Python induced modifications to the original DHCP message are also displayed in the debugging output (inserting option [51] in this case).

Python script:

```
from alc import dhcpv4
myopt = dhcpv4.getOptionList()
print "option-list =", repr(myopt)
#lease 500s 8min 20sec
dhcpv4.set(51,('\x00\x00\x01\xf4',))
```

Debug Output:

```
32826 2014/07/24 03:02:46.44 WEST MINOR: DEBUG #2001 Base Python Output
"Python Output: acg
option-list  (53, 54, 50, 82, 255)
"
32827 2014/07/24 03:02:46.44 WEST MINOR: DEBUG #2001 Base Python Result
"Python Result: acg
DHCPv4 Option 51, SET
        '\x00\x00\x01\xf4'
"
```

DHCPV6

Also the DHCPv6 Python script has some lines added to demonstate Python debugging capabilities. The new lines print assigned values to the debugging output.

DHCPv6 script

```
from alc import dhcpv6
import struct
packet = dhcpv6.get_relaymsg()
msgTop = ord(dhcpv6.msg_type)
msgBot = ord(packet.msg_type)
ia_na = packet.get_iana()
ia_pd = packet.get_iapd()
print 'ia-na = ', ia_na
print '\n'
print 'ia-pd = ', ia_pd
print '\n'
print 'msg type Top = ', msgTop
print 'msg type Bot = ', msgBot

msgType = struct.unpack('B',packet.msg_type)[0]
print "relay packet: ", msgType

# in relay request insert DHCPv6 lease times
if msgBot == 3:

        ia_na[0][1] = '\x00\x00\x07\xd0'
        ia_na[0][2] = '\x00\x00\x0b\xb8'
        ia_na[0][3][5][0][1] = '\x00\x00\x0f\xa0'
        ia_na[0][3][5][0][2] = '\x00\x00\x0f\xa0'
        packet.set_iana(ia_na)
        ia_pd[0][1] = '\x00\x00\x07\xd0'
        ia_pd[0][2] = '\x00\x00\x0b\xb8'
        ia_pd[0][3][26][0][0] = '\x00\x00\x0f\xa0'
        ia_pd[0][3][26][0][1] = '\x00\x00\x0f\xa0'
        packet.set_iapd(ia_pd)
        dhcpv6.set_relaymsg(packet)
```

Python debugging output

```
32873 2014/07/24 03:15:28.32 WEST MINOR: DEBUG #2001 Base Python Output
"Python Output: acg6

ia-na =  [['\x00\x00\x00\x00', '\x00\x00\x02X', '\x00\x00\x03\x84', {5: [[' \x01
\r\xb8\x00\x03\x00\x01\x00\x00\x00\x00\x00\x00\x00\x01', '\x00\x00\x04\xb0', '\x
00\x00\x04\xb0', {}]]}]]

ia-pd =  [['\x00\x00\x00\x00', '\x00\x00\x02X', '\x00\x00\x03\x84', {26: [['\x00
\x00\x04\xb0', '\x00\x00\x04\xb0', '6', ' \x01\r\xb8\x00\x04\x04\x00\x00\x00\x00
\x00\x00\x00\x00\x00', {}]]}]]

msg type Top =  12
msg type Bot =  3
relay packet:  3
"
```

```
32874 2014/07/24 03:15:28.32 WEST MINOR: DEBUG #2001 Base Python Result
"Python Result: acg6
DHCPv6 Option 9, SET
        '\x03\x06\x00\x00\x00\x08\x00\x02\x00\x00\x00\x01\x00\n\x00\x03\x00\x01\x00\
x00
e\x01\x03\x01\x00\x02\x00\n\x00\x03\x00\x01\xd8\x96\xff\x00\x00\x00\x00\x06\x00\
x04\x00\x03\x00\x19\x00\x03\x00(\x00\x00\x00\x00\x00\x00\x07\xd0\x00\x00\x0b\xb8
\x00\x05\x00\x18 \x01\r\xb8\x00\x03\x00\x01\x00\x00\x00\x00\x00\x00\x00\x01\x00\
x00\x0f\xa0\x00\x00\x0f\xa0\x00\x19\x00)\x00\x00\x00\x00\x00\x00\x07\xd0\x00\x00
\x0b\xb8\x00\x1a\x00\x19\x00\x00\x0f\xa0\x00\x00\x0f\xa06 \x01\r\xb8\x00\x04\x04
\x00\x00\x00\x00\x00\x00\x00\x00\x00'
"
```

Conclusion

The most common DHCP options that need to be passed by the 7x50 to the clients can be directly configured in CLI with a DHCP option specific command (such as DNS or a router option in IPv4). The DHCP option specific commands hide the complexity of the option encoding from the operator.

Less common options can be configured via a custom-option command in CLI. This scenario requires the operator to be familiar with the encoding of the option.

Similarly, RADIUS provides the means to pass the DHCP options destined to the client in the form of option specific RADIUS attributes (lease-time, etc). For less common options, two RADIUS attributes are provided: **Alc-ToClient-Dhcp-Options** and **Alc-ToClient-Dhcp6-Options**. These two attributes allow the operator to encode client destined DHCP options using hexadecimal notation. Although this process requires manual encoding it provides a very flexible way of providing options to the client.

The custom options supplied via LUDB or RADIUS are appended by the 7x50 DHCP-relay agent to any existing options that may have been already inserted by the DHCP server in the DHCP packet.

Python processing can additionally assist in DHCP message processing where the options or the parameters within the existing options can be added, removed or modified.

Routed CO

This chapter provides information about Routed Central Office (Routed CO) configurations.

Topics in this chapter include:

Applicability

This example is applicable to the 7750 SR and SR-c series as well as the 7450 ESS series in mixed mode and was tested on SR-OS 11.0.R4. Chassis mode C or higher must be used.

Summary

In the Routed Central Office (Routed CO) model, subscriber management features are implemented on a Layer 3 subscriber interface, available in a VPRN or an IES service. Compared to regular Layer 3 interfaces, a subscriber-interface supports multiple SAP’s, see later.

Customer originated traffic enters an Access Node (AN) and can be aggregated via either a Layer 2 or a Layer 3 aggregation network before being handled by a Broadband Network Gateway (BNG). Alternatively, an AN can be directly connected to the BNG.

Routed CO supports numbered, unnumbered and hybrid (combined numbered/unnumbered) subscriber interface configurations.

Enhanced Subscriber Management (ESM) is not mandatory for IPoEv4 in Routed CO, but is mandatory for PPPoE and all IPoEv6 scenarios.

The numbered and unnumbered scenarios in this example use an IES service with:

- Dual Stack IPoEv4 + IPoEv6
- Single stack PPPoEv4

General knowledge of Triple Play Service Delivery Architecture is assumed throughout this chapter. Refer to the 7x50 SR OS Triple Play Guide.

Overview

The Routed CO model offers through the subscriber and group interface construct:

- Flexible subnet management
  - Subnets can be shared across multiple access nodes.
- Support for different deployment models, for example:
  - VLAN/service model.
  - VLAN/subscriber model.
  - VLAN/service/subscriber model.
  - VLAN/access node model.
- Per group-interface load balancing in multi-chassis redundancy configurations. Redundancy is out of the scope of this example.

The components needed in the Routed CO model are depicted in Figure 215.

For the Routed CO model two interface types are needed:

- First, one or more subscriber interfaces must be created.
- Second, one or more group interfaces must be created within the subscriber interface context.

Figure 215 **Components of the Routed CO Model**

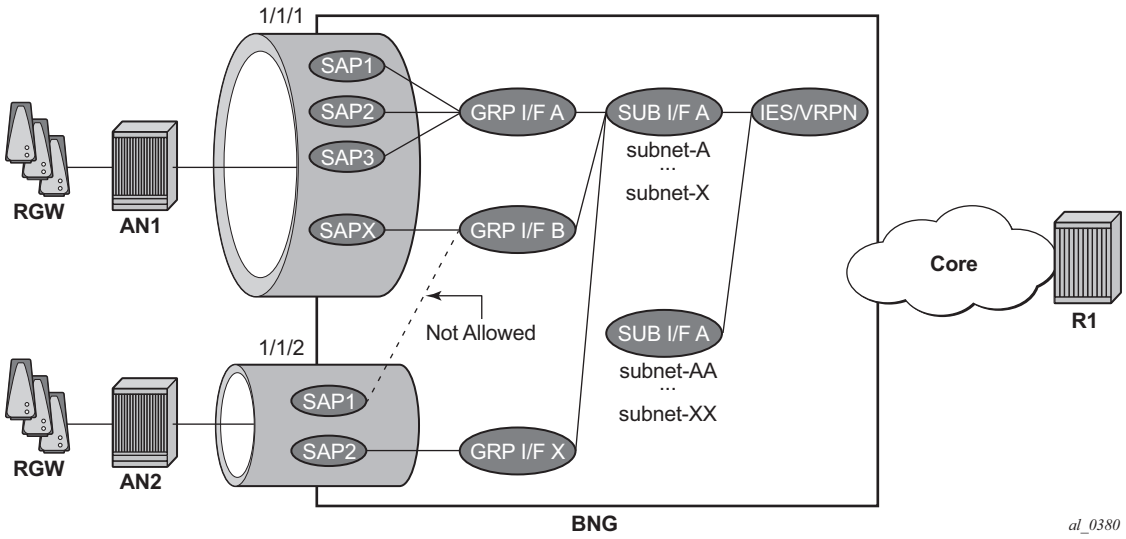

Subscriber Interface

A subscriber interface is a set of one or more group interfaces and identified by name.

A subscriber interface is created under an IES or VPRN service context, and supports up to 256 subnets (sum of IPv4 subnets and IPv6 prefixes).

Three types of subscriber interface configurations are available:

- Numbered subscriber interface.
- Unnumbered subscriber interface.
- Hybrid subscriber interface (numbered and unnumbered combined).

Subnet/Prefix Assignment

For the numbered scenario, the subscriber interface is configured with

- One or more IPv4 subnets.
- One or more IPv6 subscriber prefixes:

  - For WAN-hosts, using the DHCPv6 Identity Association for Non-temporary Addresses (IA_NA) option or Stateless Address Auto Configuration (SLAAC) and the prefix length is /64.
  - For Prefix Delegation-hosts (PD-hosts), using the DHCPv6 Identity Association for Prefix Delegation (IA_PD) option and the prefix length is defined by the Delegated Prefix Length (DPL).

This allows for subscriber-host address assignment in these subnets/prefixes only.

For the unnumbered scenario, the subscriber interface is configured with:

- IPv4:
  - No IPv4 subnets.
  - The keyword **unnumbered** plus an interface in the same routing instance (for example the system interface). The IP address of the interface referenced in the unnumbered command is used in IPCP negotiation.
- IPv6:
  - No IPv6 prefixes.
  - **allow-unmatching-prefixes**.

This allows for subscriber-host address assignment in any subnet/prefix. For IPv4, the keywords **unnumbered** and **allow-unmatching-subnets** are mutually exclusive.

For the hybrid scenario the subscriber interface is configured with:

- One or more IPv4 subnets and/or IPv6 subscriber prefixes.
- For IPv4: the keyword **allow-unmatching-subnets**.
- For IPv6: the keyword **allow-unmatching-prefixes**.

This allows for both subscriber-host address assignment within and outside of these subnets/prefixes.

Host IP Reachability

For the numbered scenario, host IP reachability requires:

- Adding the subscriber interfaces to the Interior Gateway Protocol (IGP).
- Or an export policy matching the subscriber interface subnets/prefixes.

For the unnumbered scenario, host IP reachability requires:

- An export policy matching the addresses of all individual subscriber hosts (from protocol sub-mgmt).

For the hybrid scenario, host IP reachability requires:

- An export policy matching both the subscriber interface subnets/prefixes as well as all individual subscriber hosts addresses.

Detailed examples of numbered/unnumbered/hybrid scenarios, including host IP reachability are included below.

Group Interface

A group interface is a set of one or more SAPs belonging to the same port and identified by name.

Configuration

This section covers:

- The definition of subscriber and group interfaces.
- A description of the numbered, unnumbered and hybrid scenarios.
- Options ensuring host IP reachability throughout the network.

Subscriber Interface

The configuration of the subscriber interface appears as follows.

```
configure
    service
        ies 1
            subscriber-interface "sub-int-1" create
                address 10.1.1.254/24
                address 10.1.2.254/24
                ipv6
                    delegated-prefix-len 56
                    link-local-address FE80::EA:48:FF
                    subscriber-prefixes
                        prefix 2001:DB8:101::/48 wan-host
                        prefix 2001:DB8:102::/48 pd
                        prefix 2001:DB8:F101::/48 wan-host
```

```
                prefix 2001:DB8:F102::/48 pd
            exit
        exit
    exit
    subscriber-interface "sub-int-2" create
        address 10.2.1.254/24
        address 10.2.2.254/24
        ipv6
            delegated-prefix-len 56
            link-local-address FE80::EA:48:FF
            subscriber-prefixes
                prefix 2001:DB8:201::/48 wan-host
                prefix 2001:DB8:202::/48 pd
                prefix 2001:DB8:F201::/48 wan-host
                prefix 2001:DB8:F202::/48 pd
            exit
        exit
    exit
```

Notice that once a subnet/prefix is assigned to a subscriber interface, the subnet/prefix is tied to that interface, meaning that the same subnet/prefix cannot be used on another subscriber interface or regular interface in the same routing instance. When using VPRN for the Routed CO model, overlapping subnets/prefixes are allowed when on different VPRN services.

As long as no group interfaces are configured within the subscriber interface context, the subscriber interfaces are in the operationally down state as shown in the following output. The subscriber-interfaces, sub-int-1 and sub-int-2, are operational down since no group-interfaces have been assigned at this stage.

```
*A:BNG# show router "Base" interface
===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm         Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                    PfxState
-------------------------------------------------------------------------------
sub-int-1                        Up          Down/Down   IES Sub subscriber
   10.1.1.254/24                                                 n/a
   10.1.2.254/24                                                 n/a
   2001:DB8:101::/48                                             INACCESSIBLE
   2001:DB8:102::/48                                             INACCESSIBLE
   2001:DB8:F101::/48                                            INACCESSIBLE
   2001:DB8:F102::/48                                            INACCESSIBLE
   FE80::EA:48:FF/64                                             INACCESSIBLE
sub-int-2                        Up          Down/Down   IES Sub subscriber
   10.2.1.254/24                                                 n/a
   10.2.2.254/24                                                 n/a
   2001:DB8:201::/48                                             INACCESSIBLE
   2001:DB8:202::/48                                             INACCESSIBLE
   2001:DB8:F201::/48                                            INACCESSIBLE
   2001:DB8:F202::/48                                            INACCESSIBLE
   FE80::EA:48:FF/64                                             INACCESSIBLE
system                           Up          Up/Up       Network system
   192.0.2.75/32                                                 n/a
```

```
   2001:DB8::75/128                                              PREFERRED
toDHCP-1                        Up          Up/Up       Network loopback
   10.11.11.1/32                                                 n/a
   2001:DB8::11/128                                              PREFERRED
   FE80::E84B:FFFF:FE00:0/64                                     PREFERRED
toRADIUS-1                      Up          Up/Down     Network 1/1/10
   192.168.202.75/24                                             n/a
-------------------------------------------------------------------------------
Interfaces : 5
===============================================================================
*A:BNG#
```

The corresponding IPv4 routing table looks as follows.

```
*A:BNG# show router "Base" route-table ipv4
===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
10.11.11.1/32                                   Local   Local     00h30m12s  0
       toDHCP-1                                                       0
192.0.2.75/32                                   Local   Local     00h30m12s  0
       system                                                         0
192.168.202.0/24                                Local   Local     00h29m54s  0
       toRADIUS-1                                                     0
-------------------------------------------------------------------------------
No. of Routes: 3
Flags: L = LFA nexthop available    B = BGP backup route available
       n = Number of times nexthop is repeated
===============================================================================
*A:BNG#
```

The corresponding IPv6 routing table looks as follows.

```
*A:BNG# show router "Base" route-table ipv6
===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                    Metric
-------------------------------------------------------------------------------
2001:DB8::11/128                                Local   Local     00h30m19s  0
       toDHCP-1                                                       0
2001:DB8::75/128                                Local   Local     00h30m21s  0
       system                                                         0
-------------------------------------------------------------------------------
No. of Routes: 2
Flags: L = LFA nexthop available    B = BGP backup route available
       n = Number of times nexthop is repeated
===============================================================================
*A:BNG#
```

No subscriber interface subnets/prefixes are present in the IPv4 and the IPv6 routing table as the subscriber interfaces are operational down.

Group Interface

A group interface is created under the subscriber-interface hierarchy.

```
configure
    service
        ies 1
            subscriber-interface "sub-int-1" create
                group-interface "grp-int-1-1" create
                    ipv6
                    exit
                    sap 1/1/1:111 create
                    exit
                    sap 1/1/1:112 create
                    exit
                exit
                group-interface "grp-int-1-2" create
                    ipv6
                    exit
                    sap 1/1/1:121 create
                    exit
                exit
            exit
            subscriber-interface "sub-int-2" create
                group-interface "grp-int-2-1" create
                    ipv6
                    exit
                    sap 1/1/2:211 create
                    exit
                exit
                group-interface "grp-int-2-2" create
                    ipv6
                    exit
                    sap 1/1/3:221 create
                    exit
                    sap 1/1/3:222 create
                    exit
                exit
            exit
        exit
```

Static SAPs are created manually under the group-interface context. Managed SAPs (MSAPs) are dynamically created when a trigger packet (DHCP, DHCPv6, ARP, PPPoE) is successfully authenticated, which eliminates the provisioning of static SAPs. The creation and use of capture and managed SAPs (MSAPs) is explained in the example on Managed SAPs with Routed CO.

A group interface is operationally up when at least one of its statically configured SAPs is operationally up or when no static SAPs are configured while the parameter **oper-up-while-empty** under the group-interface context is enabled. The following output shows all group interfaces are operationally up.

```
*A:BNG# show router "Base" interface ipv4
===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm         Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                    PfxState
-------------------------------------------------------------------------------
grp-int-1-1                      Up          Up/Up       IES Grp 1/1/1
grp-int-1-2                      Up          Up/Up       IES Grp 1/1/1
grp-int-2-1                      Up          Up/Up       IES Grp 1/1/2
grp-int-2-2                      Up          Up/Up       IES Grp 1/1/3
sub-int-1                        Up          Up/Up       IES Sub subscriber
   10.1.1.254/24                                                 n/a
   10.1.2.254/24                                                 n/a
sub-int-2                        Up          Up/Up       IES Sub subscriber
   10.2.1.254/24                                                 n/a
   10.2.2.254/24                                                 n/a
system                           Up          Up/Up       Network system
   192.0.2.75/32                                                 n/a
toDHCP-1                         Up          Up/Up       Network loopback
   10.11.11.1/32                                                 n/a
toRADIUS-1                       Up          Up/Down     Network 1/1/10
   192.168.202.75/24                                             n/a
-------------------------------------------------------------------------------
Interfaces : 9
===============================================================================
*A:BNG#
```

The IPv4 routing table includes the subnets configured on the subscriber-interfaces.

```
*A:BNG# show router "Base" route-table ipv4
===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.1.1.0/24                                     Local   Local     00h25m32s  0
       sub-int-1                                                       0
10.1.2.0/24                                     Local   Local     00h25m32s  0
       sub-int-1                                                       0
10.2.1.0/24                                     Local   Local     00h25m32s  0
       sub-int-2                                                       0
10.2.2.0/24                                     Local   Local     00h25m32s  0
       sub-int-2                                                       0
10.11.11.1/32                                   Local   Local     01h01m53s  0
       toDHCP-1                                                        0
192.0.2.75/32                                   Local   Local     01h01m53s  0
       system                                                          0
192.168.202.0/24                                Local   Local     01h01m36s  0
       toRADIUS-1                                                      0
```

```
-------------------------------------------------------------------------------
No. of Routes: 7
Flags: L = LFA nexthop available    B = BGP backup route available
       n = Number of times nexthop is repeated
===============================================================================
*A:BNG#
```

For IPv6, the interface table looks as follows.

```
*A:BNG# show router "Base" interface ipv6
===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm         Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                    PfxState
-------------------------------------------------------------------------------
grp-int-1-1                      Up          Up/Up       IES Grp 1/1/1
grp-int-1-2                      Up          Up/Up       IES Grp 1/1/1
grp-int-2-1                      Up          Up/Up       IES Grp 1/1/2
grp-int-2-2                      Up          Up/Up       IES Grp 1/1/3
sub-int-1                        Up          Up/Up       IES Sub subscriber
   2001:DB8:101::/48                                             PREFERRED
   2001:DB8:102::/48                                             PREFERRED
   2001:DB8:F101::/48                                            PREFERRED
   2001:DB8:F102::/48                                            PREFERRED
   FE80::EA:48:FF/64                                             PREFERRED
sub-int-2                        Up          Up/Up       IES Sub subscriber
   2001:DB8:201::/48                                             PREFERRED
   2001:DB8:202::/48                                             PREFERRED
   2001:DB8:F201::/48                                            PREFERRED
   2001:DB8:F202::/48                                            PREFERRED
   FE80::EA:48:FF/64                                             PREFERRED
system                           Up          Up/Up       Network system
   2001:DB8::75/128                                              PREFERRED
toDHCP-1                         Up          Up/Up       Network loopback
   2001:DB8::11/128                                              PREFERRED
   FE80::E84B:FFFF:FE00:0/64                                     PREFERRED
toRADIUS-1                       Up          Up/Down     Network 1/1/10
   -                                                             -
-------------------------------------------------------------------------------
Interfaces : 9
===============================================================================
*A:BNG# #
```

The IPv6 routing table includes the prefixes configured on the subscriber interfaces.

```
*A:BNG# show router "Base" route-table ipv6
===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
2001:DB8::11/128                                Local   Local     14d04h08m  0
       toDHCP-1                                                        0
```

```
2001:DB8::75/128                                Local   Local     14d04h08m  0
       system                                                          0
2001:DB8:101::/48                               Local   Local     14d03h10m  0
       sub-int-1                                                       0
2001:DB8:102::/48                               Local   Local     14d03h10m  0
       sub-int-1                                                       0
2001:DB8:201::/48                               Local   Local     14d03h08m  0
       sub-int-2                                                       0
2001:DB8:202::/48                               Local   Local     14d03h08m  0
       sub-int-2                                                       0
2001:DB8:F101::/48                              Local   Local     14d03h10m  0
       sub-int-1                                                       0
2001:DB8:F102::/48                              Local   Local     14d03h10m  0
       sub-int-1                                                       0
2001:DB8:F201::/48                              Local   Local     14d03h08m  0
       sub-int-2                                                       0
2001:DB8:F202::/48                              Local   Local     14d03h08m  0
       sub-int-2                                                       0
-------------------------------------------------------------------------------
No. of Routes: 10
Flags: L = LFA nexthop available    B = BGP backup route available
       n = Number of times nexthop is repeated
===============================================================================
*A:BNG#
```

Numbered Scenario

Figure 216 depicts the numbered scenario outlined below, including the connecting subscribers and subscriber hosts. Subscribers sub-11 and sub-44 are using PPPv4 hosts, and subscribers sub-22 and sub-33 are using dual stack DHCP hosts. Their VLANs and the MAC addresses are shown, as are the IP addresses and prefixes assigned once they are connected.

Figure 216 **Numbered Scenario For IES 1**

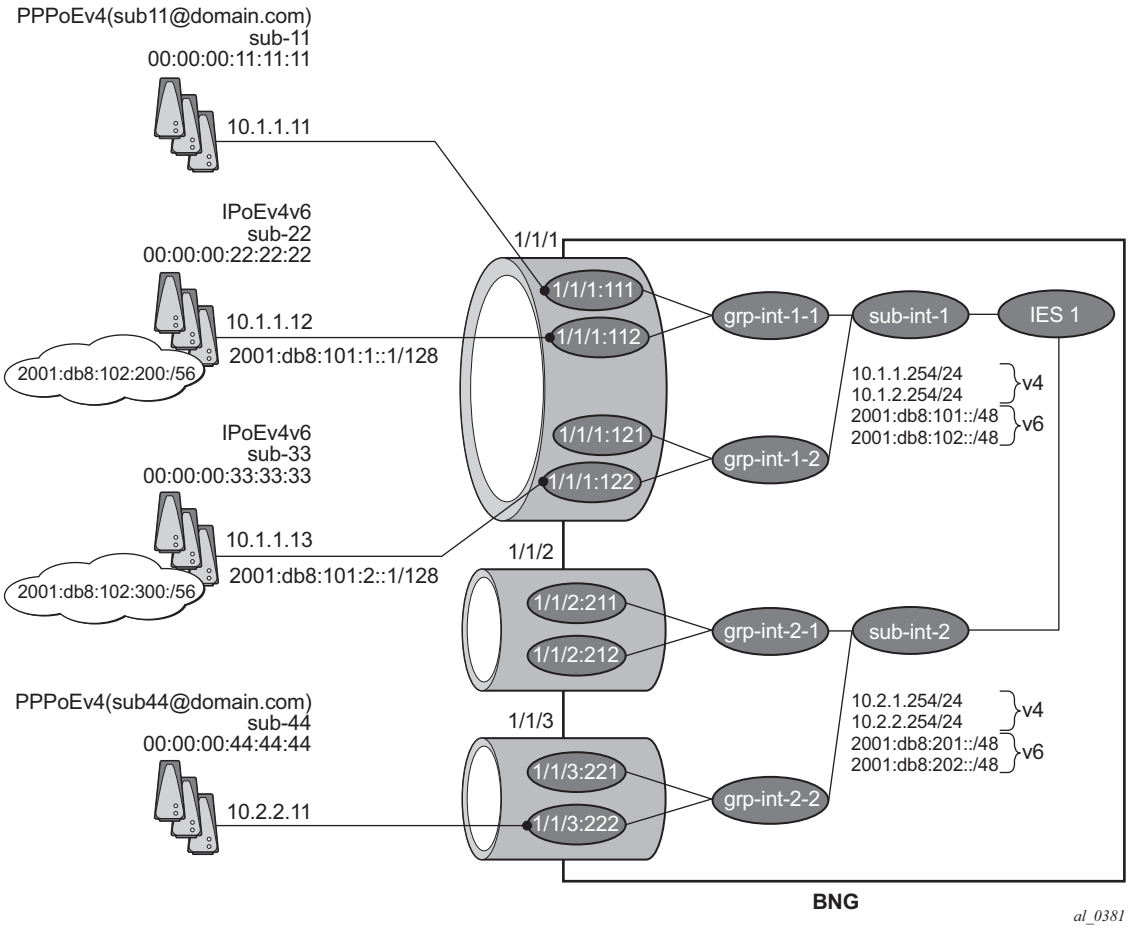

The configuration for the numbered scenario is shown below. Only the configuration items specific to the numbered scenario are shown.

In the numbered scenario the subscriber interfaces have following configuration:

- IPv4
  - Subnets.
  - **no allow-unmatching-subnets.**
  - **no unnumbered.**
- IPv6
  - A delegated prefix length.
  - subscriber prefixes.
  - **no allow-unmatching-prefixes.**

```
configure
    service
        ies 1
            subscriber-interface "sub-int-1" create
                address 10.1.1.254/24
                address 10.1.2.254/24
                ipv6
                    delegated-prefix-len 56
                    link-local-address FE80::EA:4B:FF
                    subscriber-prefixes
                        prefix 2001:DB8:101::/48 wan-host
                        prefix 2001:DB8:102::/48 pd
                    exit
                exit
                group-interface "grp-int-1-1" create
                    ipv6
                        --- snip ---
                    exit
                    arp-populate
                    dhcp
                        --- snip ---
                        lease-populate 100
                        no shutdown
                    exit
                    authentication-policy "auth-pol-1"
                    local-proxy-arp
                    sap 1/1/1:111 create
                        anti-spoof ip-mac
                        sub-sla-mgmt
                        --- snip ---
                        exit
                    exit
                    sap 1/1/1:112 create
                        anti-spoof ip-mac
                        sub-sla-mgmt
                        --- snip ---
                        exit
                    exit
                    pppoe
                        --- snip ---
                        no shutdown
                    exit
                exit
                group-interface "grp-int-1-2" create
                    ipv6
                        --- snip ---
                    exit
                    arp-populate
                    dhcp
                        --- snip ---
                        lease-populate 100
                        no shutdown
                    exit
                    authentication-policy "auth-pol-1"
                    local-proxy-arp
                    sap 1/1/1:121 create
                        anti-spoof ip-mac
                        sub-sla-mgmt
                        --- snip ---
```

```
                exit
            exit
            sap 1/1/1:122 create
                anti-spoof ip-mac
                sub-sla-mgmt
                --- snip ---
                exit
            exit
            pppoe
                --- snip ---
                no shutdown
            exit
        exit
    exit
    subscriber-interface "sub-int-2" create
        address 10.2.1.254/24
        address 10.2.2.254/24
        ipv6
            delegated-prefix-len 56
            link-local-address FE80::EA:4B:FF
            subscriber-prefixes
                prefix 2001:DB8:201::/48 wan-host
                prefix 2001:DB8:202::/48 pd
            exit
        exit
        group-interface "grp-int-2-1" create
            ipv6
                --- snip ---
            exit
            arp-populate
            dhcp
                --- snip ---
                lease-populate 100
                no shutdown
            exit
            authentication-policy "auth-pol-1"
            local-proxy-arp
            sap 1/1/2:211 create
                anti-spoof ip-mac
                sub-sla-mgmt
                --- snip ---
                exit
            exit
            sap 1/1/2:212 create
                anti-spoof ip-mac
                sub-sla-mgmt
                --- snip ---
                exit
            exit
            pppoe
                --- snip ---
                no shutdown
            exit
        exit
        group-interface "grp-int-2-2" create
            ipv6
                --- snip ---
            exit
            arp-populate
```

```
            dhcp
                --- snip ---
                lease-populate 100
                no shutdown
            exit
            authentication-policy "auth-pol-1"
            local-proxy-arp
            sap 1/1/3:221 create
                anti-spoof ip-mac
                sub-sla-mgmt
                --- snip ---
                exit
            exit
            sap 1/1/3:222 create
                anti-spoof ip-mac
                sub-sla-mgmt
                --- snip ---
                exit
            exit
            pppoe
                --- snip ---
                no shutdown
            exit
        exit
    exit
    no shutdown
```

The following parameters are mandatory for the routed CO model:

- **lease-populate** — DHCPv4 lease state population is enabled by default on a group-interface with DHCPv4 configured as **no shutdown**. The number of leases allowed on each SAP of the group-interface must be configured. By default one single DHCPv4 host is allowed on each SAP. This parameter enables the creation of an ESM host table entry for each DHCPv4 lease. For DHCPv6 the ESM host table entry creation is implicit: no CLI parameter is required.
- **arp-populate** — The ARP table is populated with dynamically learned entries from the DHCP lease state table or static entries from the static host table. The BNG does not send downstream ARPs for those managed ARP table entries.
- **local-proxy-arp** — Enables user to user traffic in a split-horizon environment. The BNG responds with its own MAC address to ARP requests targeting subnets configured on the subscriber interface. If the ARP request is targeting a host of the same subscriber on the same SAP, the ARP request is silently discarded. This prevents traffic within a single bridged home to be attracted to the BNG. Local-proxy-arp is enabled by default.
- **anti-spoof** — Checks the source MAC and/or source IP of the upstream subscriber traffic. This parameter is configured at the SAP level with values **ip-mac** (default), **ip** or **nh-mac**. With ESM enabled, anti-spoof must include the source mac (values **ip-mac** or **nh-mac**).

Optional settings are:

- **description** — Can be used to assign a descriptive text to the item and used for administrative reasons.
- **delayed-enable** — To be used in redundant configurations. It is expressed in seconds and defines the additional time the BNG waits before the interface is enabled.

Verification

The interfaces on the BNG are listed using following command. Notice that all subscriber and group interfaces are operational up for IPv4 and IPv6.

```
A:BNG# show router "Base" interface
===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm          Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                    PfxState
-------------------------------------------------------------------------------
grp-int-1-1                      Up           Up/Up       IES Grp 1/1/1
grp-int-1-2                      Up           Up/Up       IES Grp 1/1/1
grp-int-2-1                      Up           Up/Up       IES Grp 1/1/2
grp-int-2-2                      Up           Up/Up       IES Grp 1/1/3
sub-int-1                        Up           Up/Up       IES Sub subscriber
   10.1.1.254/24                                                 n/a
   10.1.2.254/24                                                 n/a
   2001:DB8:101::/48                                             PREFERRED
   2001:DB8:102::/48                                             PREFERRED
   FE80::EA:48:FF/64                                             PREFERRED
sub-int-2                        Up           Up/Up       IES Sub subscriber
   10.2.1.254/24                                                 n/a
   10.2.2.254/24                                                 n/a
   2001:DB8:201::/48                                             PREFERRED
   2001:DB8:202::/48                                             PREFERRED
   FE80::EA:48:FF/64                                             PREFERRED
system                           Up           Up/Up       Network system
   192.0.2.75/32                                                 n/a
   2001:DB8::75/128                                              PREFERRED
toDHCP-1                         Up           Up/Up       Network loopback
   10.11.11.1/32                                                 n/a
   2001:DB8::11/128                                              PREFERRED
   FE80::E84B:FFFF:FE00:0/64                                     PREFERRED
toR1                             Up           Up/Up       Network 1/1/12
   192.168.12.1/24                                               n/a
   2001:DEAD::1/64                                               PREFERRED
   FE80::E84B:FFFF:FE00:0/64                                     PREFERRED
toRADIUS-1                       Up           Up/Down     Network 1/1/10
   192.168.202.75/24                                             n/a
-------------------------------------------------------------------------------
Interfaces : 10
===============================================================================
A:BNG#
```

Successfully created hosts have forwarding state Fwding. Hosts not in the Fwding state cannot forward any data.

```
A:BNG# show service id 1 subscriber-hosts
===============================================================
Subscriber Host table
===============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID Origin       Fwding State
---------------------------------------------------------------
1/1/1:111              sub-11
  10.1.1.11
    00:00:00:11:11:11    1         IPCP         Fwding
1/1/1:112              sub-22
  10.1.1.12
    00:00:00:22:22:22    N/A       DHCP         Fwding
1/1/1:112              sub-22
  2001:DB8:101:1::1/128
    00:00:00:22:22:22    N/A       IPoE-DHCP6   Fwding
1/1/1:112              sub-22
  2001:DB8:102:200::/56
    00:00:00:22:22:22    N/A       IPoE-DHCP6   Fwding
1/1/1:122              sub-33
  10.1.1.13
    00:00:00:33:33:33    N/A       DHCP         Fwding
1/1/1:122              sub-33
  2001:DB8:101:2::1/128
    00:00:00:33:33:33    N/A       IPoE-DHCP6   Fwding
1/1/1:122              sub-33
  2001:DB8:102:300::/56
    00:00:00:33:33:33    N/A       IPoE-DHCP6   Fwding
1/1/3:222              sub-44
  10.2.2.11
    00:00:00:44:44:44    1         IPCP         Fwding
---------------------------------------------------------------
Number of subscriber hosts : 8
===============================================================
A:BNG#
```

The list of active subscribers can be displayed as follows.

```
A:BNG# show service active-subscribers
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber sub-11 (sub-prof-1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:1/1/1:111 - sla:sla-prof-1
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.1.1.11
                00:00:00:11:11:11 1         IPCP
```

```
-------------------------------------------------------------------------------
Subscriber sub-22 (sub-prof-1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:1/1/1:112 - sla:sla-prof-1
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.1.1.12
                00:00:00:22:22:22 N/A       DHCP
2001:DB8:101:1::1/128
                00:00:00:22:22:22 N/A       IPoE-DHCP6
2001:DB8:102:200::/56
                00:00:00:22:22:22 N/A       IPoE-DHCP6
-------------------------------------------------------------------------------
Subscriber sub-33 (sub-prof-1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:1/1/1:122 - sla:sla-prof-1
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.1.1.13
                00:00:00:33:33:33 N/A       DHCP
2001:DB8:101:2::1/128
                00:00:00:33:33:33 N/A       IPoE-DHCP6
2001:DB8:102:300::/56
                00:00:00:33:33:33 N/A       IPoE-DHCP6
-------------------------------------------------------------------------------
Subscriber sub-44 (sub-prof-1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:1/1/3:222 - sla:sla-prof-1
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.2.2.11
                00:00:00:44:44:44 1         IPCP
-------------------------------------------------------------------------------
Number of active subscribers : 4
-------------------------------------------------------------------------------
A:BNG#
```

Manually cross-referencing the SAPs from this output with the actual configuration shows the following for IPv4, and is depicted in Figure 216.

- Sub-11 and sub-22 are connected to the same subscriber and group interface (sub-int-1 and grp-int-1-1) but via different SAPs (1/1/1:111 and 1/1/1:112) and are sharing the same IPv4 subnet.

- Sub-33 is also connected to the same subscriber interface (sub-int-1) but via a different group-interface (grp-int-1-2). Sub-33 shares the same IPv4 subnet as sub-11 and sub-12, showing that the same subnet is shared across multiple group-interfaces.
- Sub-44 is connected to a different subscriber and group interface, and does not share a subnet with the other subscribers.

An alternative way to find where, for example, subscriber sub-33 is connected is shown below.

```
*A:BNG# show service active-subscribers subscriber "sub-33" detail
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber sub-11 (sub-prof-1)
-------------------------------------------------------------------------------
I. Sched. Policy : N/A

   ---  snip ---

Oper-Rate-Limit  : Maximum
* indicates that the corresponding row element may have been truncated.
-------------------------------------------------------------------------------
(1) SLA Profile Instance
    - sap:1/1/1:112 (IES 1 - grp-int-1-2)
    - sla:sla-prof-1
-------------------------------------------------------------------------------
Description          : (Not Specified)

   ---  snip ---
```

An alternative to find where, for example, IP address 10.1.1.13 is connected is shown below.

```
*A:BNG# show service id 1 dhcp lease-state ip-address 10.1.1.13 detail
===============================================================================
DHCP lease states for service 1
===============================================================================
Service ID           : 1
IP Address           : 10.1.1.13
Client HW Address    : 00:00:00:33:33:33
Subscriber-interface : sub-int-1
Group-interface      : grp-int-1-2
SAP                  : 1/1/1:122
--- snip ---

Sub-Ident            : "sub-33"
Sub-Profile-String   : "sub-prof-1"
SLA-Profile-String   : "sla-prof-1"
App-Profile-String   : ""

--- snip ---
```

```
DHCP Server Addr     : 10.11.11.1
Radius User-Name     : "00:00:00:33:33:33"
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
*A:BNG#
```

For IPv6, the situation is as follows:

- Sub-22 and sub-33 are connected to the same subscriber interface (sub-int-1) but to different group interfaces. Both subscribers share the same IPv6 prefix for prefix-delegation (PD) and wan-host.

With these subscriber hosts connected, the IPv4 routing table (RIB) for the base router looks as follows.

```
*A:BNG# show router "Base" route-table ipv4
===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.1.1.0/24                                     Local   Local     02h25m15s  0
       sub-int-1                                                        0
10.1.1.11/32                                    Remote  Sub Mgmt  02h25m10s  0
       [grp-int-1-1]                                                    0
10.1.1.12/32                                    Remote  Sub Mgmt  00h49m52s  0
       [grp-int-1-1]                                                    0
10.1.1.13/32                                    Remote  Sub Mgmt  00h47m40s  0
       [grp-int-1-2]                                                    0
10.1.2.0/24                                     Local   Local     02h25m15s  0
       sub-int-1                                                        0
10.2.1.0/24                                     Local   Local     02h25m15s  0
       sub-int-2                                                        0
10.2.2.0/24                                     Local   Local     02h25m15s  0
       sub-int-2                                                        0
10.2.2.11/32                                    Remote  Sub Mgmt  02h25m10s  0
       [grp-int-2-2]                                                    0
10.11.11.1/32                                   Local   Local     02h25m33s  0
       toDHCP-1                                                         0
192.0.2.75/32                                   Local   Local     02h25m33s  0
       system                                                           0
192.0.2.76/32                                   Remote  ISIS      02h24m43s  15
       192.168.12.2                                                     10
192.168.12.0/24                                 Local   Local     02h25m15s  0
       toR1                                                             0
192.168.202.0/24                                Local   Local     02h25m15s  0
       toRADIUS-1                                                       0
-------------------------------------------------------------------------------
No. of Routes: 13
Flags: L = LFA nexthop available    B = BGP backup route available
       n = Number of times nexthop is repeated
===============================================================================
*A:BNG#
```

The IPv6 routing table (RIB) for the base router displays as follows.

```
*A:BNG# show router "Base" route-table ipv6
===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age         Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
2001:DB8::11/128                                Local   Local     02h27m23s   0
       toDHCP-1                                                         0
2001:DB8::75/128                                Local   Local     02h27m24s   0
       system                                                           0
2001:DB8::76/128                                Remote  ISIS      02h26m32s   15
       FE80::E84C:FFFF:FE00:0-"toR1"                                    10
2001:DB8:101::/48                               Local   Local     02h27m06s   0
       sub-int-1                                                        0
2001:DB8:101:1::1/128                           Remote  Sub Mgmt  01h13m43s   0
       [grp-int-1-1]                                                    0
2001:DB8:101:2::1/128                           Remote  Sub Mgmt  01h13m25s   0
       [grp-int-1-2]                                                    0
2001:DB8:102::/48                               Local   Local     02h27m06s   0
       sub-int-1                                                        0
2001:DB8:102:200::/56                           Remote  Sub Mgmt  01h13m43s   0
       [grp-int-1-1]                                                    0
2001:DB8:102:300::/56                           Remote  Sub Mgmt  01h13m25s   0
       [grp-int-1-2]                                                    0
2001:DB8:201::/48                               Local   Local     02h27m06s   0
       sub-int-2                                                        0
2001:DB8:202::/48                               Local   Local     02h27m06s   0
       sub-int-2                                                        0
2001:DEAD::/64                                  Local   Local     02h27m05s   0
       toR1                                                             0
-------------------------------------------------------------------------------
No. of Routes: 12
Flags: L = LFA nexthop available    B = BGP backup route available
       n = Number of times nexthop is repeated
===============================================================================
*A:BNG#
```

The corresponding IPv4 FIB on card 1 looks as follows.

```
*A:BNG# show router "Base" fib 1 ipv4
===============================================================================
FIB Display
===============================================================================
Prefix                                                      Protocol
    NextHop
-------------------------------------------------------------------------------
10.1.1.0/24                                                 LOCAL
    10.1.1.0 (sub-int-1)
10.1.2.0/24                                                 LOCAL
    10.1.2.0 (sub-int-1)
10.2.1.0/24                                                 LOCAL
    10.2.1.0 (sub-int-2)
10.2.2.0/24                                                 LOCAL
    10.2.2.0 (sub-int-2)
```

```
10.11.11.1/32                                                       LOCAL
    10.11.11.1 (toDHCP-1)
192.0.2.75/32                                                       LOCAL
    192.0.2.75 (system)
192.0.2.76/32                                                       ISIS
    192.168.12.2 (toR1)
192.168.12.0/24                                                     LOCAL
    192.168.12.0 (toR1)
192.168.202.0/24                                                    LOCAL
    192.168.202.0 (toRADIUS-1)
-------------------------------------------------------------------------------
Total Entries : 9
-------------------------------------------------------------------------------
===============================================================================
*A:BNG#
```

The corresponding IPv6 FIB on card 1 is as follows.

```
*A:BNG# show router "Base" fib 1 ipv6
===============================================================================
FIB Display
===============================================================================
Prefix                                                              Protocol
    NextHop
-------------------------------------------------------------------------------
2001:DB8::11/128                                                    LOCAL
    2001:DB8::11 (toDHCP-1)
2001:DB8::75/128                                                    LOCAL
    2001:DB8::75 (system)
2001:DB8::76/128                                                    ISIS
    FE80::E84C:FFFF:FE00:0 (toR1)
2001:DB8:101::/48                                                   LOCAL
    2001:DB8:101:: (sub-int-1)
2001:DB8:102::/48                                                   LOCAL
    2001:DB8:102:: (sub-int-1)
2001:DB8:201::/48                                                   LOCAL
    2001:DB8:201:: (sub-int-2)
2001:DB8:202::/48                                                   LOCAL
    2001:DB8:202:: (sub-int-2)
2001:DEAD::/64                                                      LOCAL
    2001:DEAD:: (toR1)
-------------------------------------------------------------------------------
Total Entries : 8
-------------------------------------------------------------------------------
===============================================================================
*A:BNG#
```

The addresses of the individual subscriber hosts show up in the RIB but they do not show up in the FIB.

- /32 for IPv4-hosts.
- /DPL (Delegated Prefix Length) for IPv6 DP hosts, /56 in this example.
- /128 or /64 for IPv6 wan host.

Downstream traffic is forwarded based on a subscriber host table lookup. For specific network designs, subscriber host IPv4 addresses can optionally be included in the FIB with the populate-host-routes statement added to the subnet configuration. This is out of scope of this example.

Unnumbered Scenario

Figure 217 depicts the unnumbered scenario outlined below, including the connecting subscribers and subscriber hosts. Sub-11 and sub-44 are using single stack PPPoEv4 hosts, and sub-22 and sub-33 are using dual stack DHCP hosts. Their VLANs and the MAC addresses are shown, as are the IP addresses and prefixes assigned once they are connected.

Figure 217 **Unnumbered Scenario for IES 1**

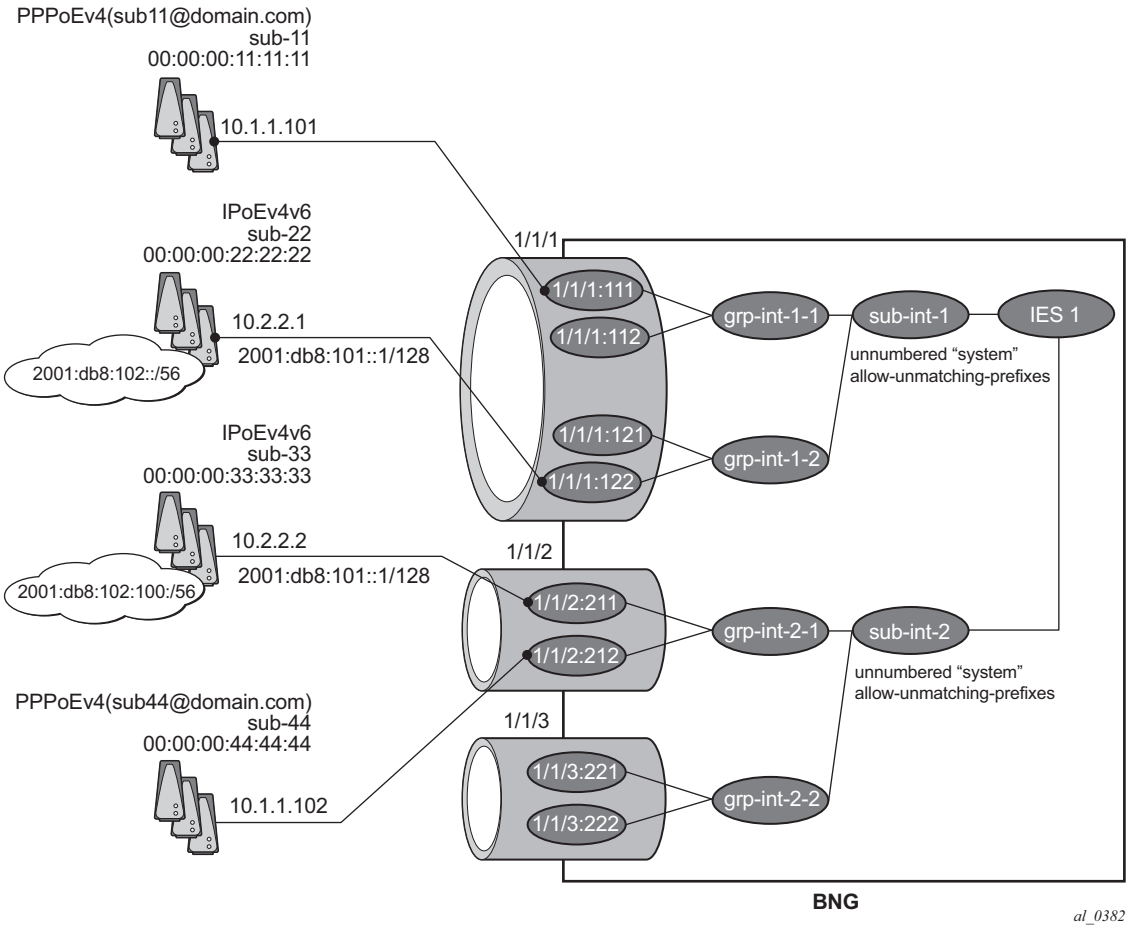

The configuration for the unnumbered scenario is show below. Only the configuration items specific to the unnumbered scenario are shown.

In the unnumbered scenario the subscriber interfaces have following properties:

- IPv4:
  - No subnets configured.
  - **unnumbered**, with an IPv4 interface or an IPv4 address used for IPCP negotiation.
  - **no allow-unmatching-subnets.**
- IPv6:
  - No subscriber prefixes configured.
  - **allow-unmatching-prefixes.**

```
configure
    service
        ies 1
            subscriber-interface "sub-int-1" create
                unnumbered "system"
                ipv6
                    delegated-prefix-len 56
                    allow-unmatching-prefixes
                    link-local-address FE80::EA:4B:FF
                exit
                group-interface "grp-int-1-1" create
                    ipv6
                        --- snip ---
                    exit
                    arp-populate
                    dhcp
                        --- snip ---
                        lease-populate 100
                        no shutdown
                    exit
                    authentication-policy "auth-pol-1"
                    sap 1/1/1:111 create
                        anti-spoof ip-mac
                        sub-sla-mgmt
                            --- snip ---
                        exit
                    exit
                    sap 1/1/1:112 create
                        anti-spoof ip-mac
                        sub-sla-mgmt
                            --- snip ---
                        exit
                    exit
                    pppoe
                        --- snip ---
                        no shutdown
                    exit
                exit
                group-interface "grp-int-1-2" create
```

```
            ipv6
                --- snip ---
            exit
            arp-populate
            dhcp
                --- snip ---
                lease-populate 100
                no shutdown
            exit
            authentication-policy "auth-pol-1"
            sap 1/1/1:121 create
                anti-spoof ip-mac
                sub-sla-mgmt
                    --- snip ---
                exit
            exit
            sap 1/1/1:122 create
                anti-spoof ip-mac
                sub-sla-mgmt
                    --- snip ---
                exit
            exit
            pppoe
                --- snip ---
                no shutdown
            exit
        exit
    exit
    subscriber-interface "sub-int-2" create
        unnumbered "system"
        ipv6
            delegated-prefix-len 56
            allow-unmatching-prefixes
            link-local-address FE80::EA:4B:FF
        exit
        group-interface "grp-int-2-1" create
            ipv6
                --- snip ---
            exit
            arp-populate
            dhcp
                --- snip ---
                lease-populate 100
                no shutdown
            exit
            authentication-policy "auth-pol-1"
            sap 1/1/2:211 create
                anti-spoof ip-mac
                sub-sla-mgmt
                    --- snip ---
                exit
            exit
            sap 1/1/2:212 create
                anti-spoof ip-mac
                sub-sla-mgmt
                    --- snip ---
                exit
            exit
            pppoe
```

```
                --- snip ---
                no shutdown
            exit
        exit
        group-interface "grp-int-2-2" create
            ipv6
                --- snip ---
            exit
            arp-populate
            dhcp
                --- snip ---
                lease-populate 100
                no shutdown
            exit
            authentication-policy "auth-pol-1"
            sap 1/1/3:221 create
                anti-spoof ip-mac
                sub-sla-mgmt
                    --- snip ---
                exit
            exit
            sap 1/1/3:222 create
                sub-sla-mgmt
                anti-spoof ip-mac
                sub-sla-mgmt
                    --- snip ---
                exit
                exit
            exit
            pppoe
                --- snip ---
                no shutdown
            exit
        exit
    exit
    no shutdown
```

The same mandatory and optional settings as for the numbered scenario apply.

Verification

The interfaces on the BNG are listed using following command. Notice that all subscriber and group interfaces are operational up for IPv4 and IPv6.

```
A:BNG# show router "Base" interface
===============================================================================
Interface Table (Router: Base)
===============================================================================
Interface-Name                   Adm         Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                    PfxState
-------------------------------------------------------------------------------
grp-int-1-1                      Up          Up/Up       IES Grp 1/1/1
grp-int-2-2                      Up          Up/Up       IES Grp 1/1/1
```

```
grp-int-2-1                         Up          Up/Up       IES Grp 1/1/2
grp-int-2-2                         Up          Up/Up       IES Grp 1/1/3
lb-pool4-1                          Up          Up/Down     Network loopback
   10.1.1.254/24                                            n/a
lb-pool4-2                          Up          Up/Down     Network loopback
   10.1.2.254/24                                            n/a
lb-pool4-3                          Up          Up/Down     Network loopback
   10.2.1.254/24                                            n/a
lb-pool4-4                          Up          Up/Down     Network loopback
   10.2.2.254/24                                            n/a
sub-int-1                           Up          Up/Up       IES Sub subscriber
   Unnumbered If[system]                                    n/a
   FE80::EA:4B:FF/64                                        PREFERRED
sub-int-2                           Up          Up/Up       IES Sub subscriber
   Unnumbered If[system]                                    n/a
   FE80::EA:4B:FF/64                                        PREFERRED
system                              Up          Up/Up       Network system
   192.0.2.75/32                                            n/a
   2001:DB8::75/128                                         PREFERRED
toDHCP-1                            Up          Up/Up       Network loopback
   10.11.11.1/32                                            n/a
   2001:DB8::11/128                                         PREFERRED
   FE80::E84B:FFFF:FE00:0/64                                PREFERRED
toR1                                Up          Up/Down     Network 1/1/12
   192.168.12.1/24                                          n/a
toRADIUS-1                          Up          Up/Down     Network 1/1/10
   192.168.202.75/24                                        n/a
-------------------------------------------------------------------------------
Interfaces : 14
===============================================================================
A:BNG#
```

Successfully created hosts have forwarding state Fwding. Hosts not in the Fwding state cannot forward any data.

```
*A:BNG# show service id 1 subscriber-hosts
===============================================================
Subscriber Host table
===============================================================
Sap                     Subscriber
  IP Address
    MAC Address           PPPoE-SID Origin       Fwding State
---------------------------------------------------------------
1/1/1:111               sub-11
  10.1.1.101
    00:00:00:11:11:11     1         IPCP         Fwding
1/1/1:122               sub-22
  10.2.2.1
    00:00:00:22:22:22     N/A       DHCP         Fwding
1/1/1:122               sub-22
  2001:DB8:101::1/128
    00:00:00:22:22:22     N/A       IPoE-DHCP6   Fwding
1/1/1:122               sub-22
  2001:DB8:102::/56
    00:00:00:22:22:22     N/A       IPoE-DHCP6   Fwding
1/1/2:211               sub-33
  10.2.2.2
    00:00:00:33:33:33     N/A       DHCP         Fwding
```

```
1/1/2:211              sub-33
  2001:DB8:101:1::1/128
    00:00:00:33:33:33    N/A       IPoE-DHCP6   Fwding
1/1/2:211              sub-33
  2001:DB8:102:100::/56
    00:00:00:33:33:33    N/A       IPoE-DHCP6   Fwding
1/1/2:212              sub-44
  10.1.1.102
    00:00:00:44:44:44    1         IPCP         Fwding
-------------------------------------------------------------
Number of subscriber hosts : 8
=============================================================
*A:BNG#
```

A variant of the show service active-subscribers command shows the subscriber hierarchy.

```
*A:BNG# show service active-subscribers hierarchy
===============================================================================
Active Subscriber hierarchy
===============================================================================
-- sub-11 (sub-prof-1)
   |
   |-- sap:1/1/1:111 - sla:sla-prof-1
   |   |
   |   |-- 10.1.1.101
   |   |   00:00:00:11:11:11 - 1 (IPCP)
   |   |

-- sub-22 (sub-prof-1)
   |
   |-- sap:1/1/1:122 - sla:sla-prof-1
   |   |
   |   |-- 10.2.2.1
   |   |   00:00:00:22:22:22 - N/A (DHCP)
   |   |
   |   |-- 2001:DB8:101::1/128
   |   |   00:00:00:22:22:22 - N/A (IPoE-DHCP6)
   |   |
   |   |-- 2001:DB8:102::/56
   |   |   00:00:00:22:22:22 - N/A (IPoE-DHCP6)
   |   |

-- sub-33 (sub-prof-1)
   |
   |-- sap:1/1/2:211 - sla:sla-prof-1
   |   |
   |   |-- 10.2.2.2
   |   |   00:00:00:33:33:33 - N/A (DHCP)
   |   |
   |   |-- 2001:DB8:101:1::1/128
   |   |   00:00:00:33:33:33 - N/A (IPoE-DHCP6)
   |   |
   |   |-- 2001:DB8:102:100::/56
   |   |   00:00:00:33:33:33 - N/A (IPoE-DHCP6)
   |   |
```

```
-- sub-44 (sub-prof-1)
   |
   |-- sap:1/1/2:212 - sla:sla-prof-1
   |   |
   |   |-- 10.1.1.102
   |   |   00:00:00:44:44:44 - 1 (IPCP)
   |   |
===============================================================================
*A:BNG#
```

Manually cross-referencing the SAPs from this output with the actual configuration shows the following for IPv4, and is represented in Figure 217.

- Sub-11 and sub-44 share the same IPv4 subnet even though they are connected to different subscriber interfaces.
- Sub-22 and sub-33 share the same subnet even though they are connected to different subscriber interfaces.

For IPv6 the situation is as follows:

- Sub-22 and sub-33 are in different subscriber interfaces and do not share IPv6 prefixes in this example.

With these subscriber hosts are connected, the IPv4 RIB for the base router looks as follows.

```
A:BNG# show router "Base" route-table ipv4
===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.1.1.0/24                                     Local   Local     00h49m42s  0
       lb-pool4-1                                                     0
10.1.1.101/32                                   Remote  Sub Mgmt  00h23m24s  0
       [grp-int-1-1]                                                  0
10.1.1.102/32                                   Remote  Sub Mgmt  00h02m32s  0
       [grp-int-2-1]                                                  0
10.1.2.0/24                                     Local   Local     00h49m42s  0
       lb-pool4-2                                                     0
10.2.1.0/24                                     Local   Local     00h49m42s  0
       lb-pool4-3                                                     0
10.2.2.0/24                                     Local   Local     00h49m42s  0
       lb-pool4-4                                                     0
10.2.2.1/32                                     Remote  Sub Mgmt  00h27m18s  0
       [grp-int-1-2]                                                  0
10.2.2.2/32                                     Remote  Sub Mgmt  00h27m10s  0
       [grp-int-2-1]                                                  0
10.11.11.1/32                                   Local   Local     00h49m42s  0
       toDHCP-1                                                       0
192.0.2.75/32                                   Local   Local     00h49m42s  0
       system                                                         0
192.0.2.76/32                                   Remote  ISIS      00h41m48s  15
```

```
       192.168.12.2                                                    10
192.168.12.0/24                          Local   Local     00h49m24s  0
       toR1                                                            0
192.168.202.0/24                         Local   Local     00h49m24s  0
       toRADIUS-1                                                      0
-------------------------------------------------------------------------------
No. of Routes: 13
Flags: L = LFA nexthop available    B = BGP backup route available
       n = Number of times nexthop is repeated
===============================================================================
A:BNG#
```

The IPv6 RIB for the base router looks as follows.

```
.A:BNG# show router "Base" route-table ipv6
===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                       Type    Proto     Age        Pref
      Next Hop[Interface Name]                                 Metric
-------------------------------------------------------------------------------
2001:DB8::11/128                         Local   Local     01h03m27s  0
       toDHCP-1                                                        0
2001:DB8::75/128                         Local   Local     00h06m34s  0
       system                                                          0
2001:DB8:101::1/128                      Remote  Sub Mgmt  00h36m12s  0
       [grp-int-1-2]                                                   0
2001:DB8:101:1::1/128                    Remote  Sub Mgmt  00h35m58s  0
       [grp-int-2-1]                                                   0
2001:DB8:102::/56                        Remote  Sub Mgmt  00h36m12s  0
       [grp-int-1-2]                                                   0
2001:DB8:102:100::/56                    Remote  Sub Mgmt  00h35m58s  0
       [grp-int-2-1]                                                   0
-------------------------------------------------------------------------------
No. of Routes: 6
Flags: L = LFA nexthop available    B = BGP backup route available
       n = Number of times nexthop is repeated
===============================================================================
A:BNG# #
```

The corresponding IPv4 FIB on card 1 looks as follows.

```
A:BNG# show router "Base" fib 1 ipv4
===============================================================================
FIB Display
===============================================================================
Prefix                                                          Protocol
    NextHop
-------------------------------------------------------------------------------
10.1.1.0/24                                                     LOCAL
    10.1.1.0 (lb-pool4-1)
10.1.1.101/32                                                   LOCAL
    10.1.1.101 (sub-int-1)
10.1.1.102/32                                                   LOCAL
    10.1.1.102 (sub-int-2)
```

```
10.1.2.0/24                                                         LOCAL
    10.1.2.0 (lb-pool4-2)
10.2.1.0/24                                                         LOCAL
    10.2.1.0 (lb-pool4-3)
10.2.2.0/24                                                         LOCAL
    10.2.2.0 (lb-pool4-4)
10.2.2.1/32                                                         LOCAL
    10.2.2.1 (sub-int-1)
10.2.2.2/32                                                         LOCAL
    10.2.2.2 (sub-int-2)
10.11.11.1/32                                                       LOCAL
    10.11.11.1 (toDHCP-1)
192.0.2.75/32                                                       LOCAL
    192.0.2.75 (system)
192.0.2.76/32                                                       ISIS
    192.168.12.2 (toR1)
192.168.12.0/24                                                     LOCAL
    192.168.12.0 (toR1)
192.168.202.0/24                                                    LOCAL
    192.168.202.0 (toRADIUS-1)
-------------------------------------------------------------------------------
Total Entries : 13
-------------------------------------------------------------------------------
===============================================================================
A:BNG#
```

The corresponding IPv6 FIB on card 1 looks as follows:

```
A:BNG# show router "Base" fib 1 ipv6
===============================================================================
FIB Display
===============================================================================
Prefix                                                              Protocol
    NextHop
-------------------------------------------------------------------------------
2001:DB8::11/128                                                    LOCAL
    2001:DB8::11 (toDHCP-1)
2001:DB8::75/128                                                    LOCAL
    2001:DB8::75 (system)
2001:DB8:101::1/128                                                 LOCAL
    2001:DB8:101::1 (sub-int-1)
2001:DB8:101:1::1/128                                               LOCAL
    2001:DB8:101:1::1 (sub-int-2)
2001:DB8:102::/56                                                   LOCAL
    2001:DB8:102:: (sub-int-1)
2001:DB8:102:100::/56                                               LOCAL
    2001:DB8:102:100:: (sub-int-2)
-------------------------------------------------------------------------------
Total Entries : 6
-------------------------------------------------------------------------------
===============================================================================
A:BNG#
```

The addresses of the individual subscriber hosts appear in the RIB and the FIB, which is the main difference with the numbered model. The forwarding plane here needs the individual addresses to forward the traffic towards the individual subscriber hosts.

Hybrid Scenario

An alternative to the scenarios described above does exist in the form of a mixed numbered/unnumbered (hybrid) scenario as depicted in Figure 218.

The subscriber interface is configured with

- One or more IPv4 subnets and/or IPv6 subscriber prefixes.
- For IPv4: the keyword **allow-unmatching-subnets**.
- For IPv6: the keyword **allow-unmatching-prefixes**.

No explicit configuration is shown as it is a mix of the numbered and the unnumbered scenario described above, and as such the behavior is mixed.

Figure 218 **Hybrid Configuration**

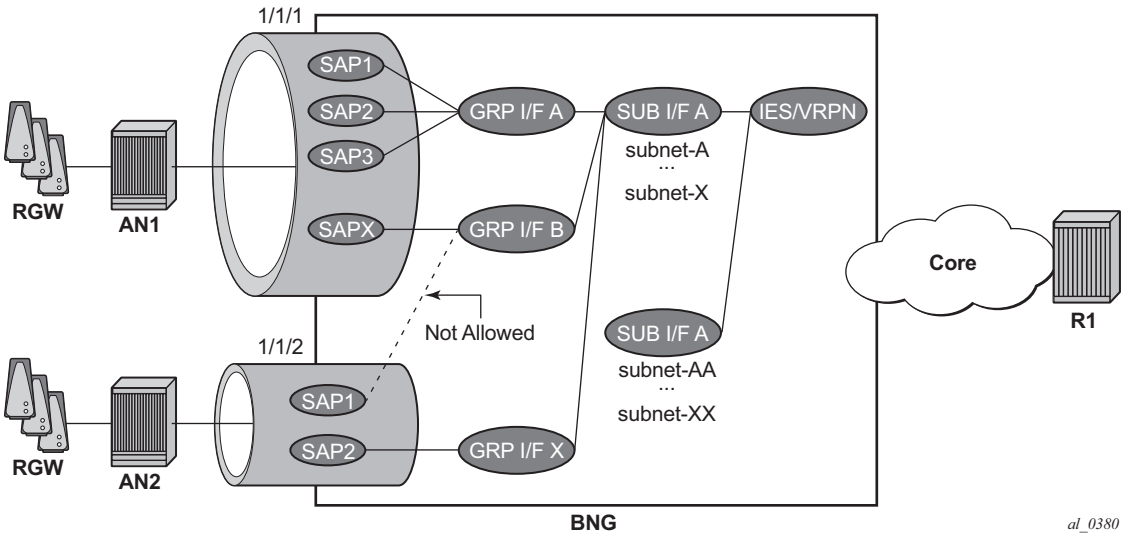

Host IP Reachability

To ensure reachability to the individual subscriber hosts, the subnets and prefixes of the subscriber interfaces/subscriber hosts need to be distributed to other routers in the network.

Three options are available:

- Without an export policy.
- With an export policy using, for example, from protocol direct.
- With an export policy using, for example, from protocol sub-mgmt.

Option 1 – No Export Policy

The key properties for the first option are:

- Subscriber interface subnets and prefixes are distributed into the network by adding the subscriber interfaces as passive interfaces to the routing protocol.
- It is used in combination with IGP based route distribution.
- It works with the numbered model only.

In this option the BNG uses IS-IS as IGP and no export policy is needed.

```
configure
    router
       isis
            area-id 48.0001
            multi-topology
                ipv6-unicast
            exit
            interface "system"
                no shutdown
            exit
            interface "sub-int-1"
                passive
                no shutdown
            exit
            interface "sub-int-2"
                passive
                no shutdown
            exit
            interface "toR1"
                interface-type point-to-point
                no shutdown
            exit
            no shutdown
        exit
```

The corresponding IPv4 RIB on router R1 (from Figure 215) lists the subscriber-interface subnets, not the individual subscriber host addresses.

```
A:R1# show router "Base" route-table ipv4
===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.1.1.0/24                                     Remote  ISIS      00h14m11s  15
       192.168.12.1                                                    20
10.1.2.0/24                                     Remote  ISIS      00h14m11s  15
       192.168.12.1                                                    20
10.2.1.0/24                                     Remote  ISIS      00h14m05s  15
       192.168.12.1                                                    20
10.2.2.0/24                                     Remote  ISIS      00h14m05s  15
       192.168.12.1                                                    20
192.0.2.75/32                                   Remote  ISIS      00h14m17s  15
       192.168.12.1                                                    10
192.0.2.76/32                                   Local   Local     62d21h22m  0
       system                                                          0
192.168.12.0/24                                 Local   Local     05d04h43m  0
       toBNG                                                           0
-------------------------------------------------------------------------------
No. of Routes: 7
Flags: L = LFA nexthop available    B = BGP backup route available
       n = Number of times nexthop is repeated
===============================================================================
A:R1#
```

The corresponding IPv6 RIB on router R1 lists the subscriber-interface prefixes, not the individual subscriber host addresses/prefixes.

```
A:R1# show router "Base" route-table ipv6
===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
2001:DB8::75/128                                Remote  ISIS      00h04m25s  15
       FE80::E84B:FFFF:FE00:0-"toBNG"                                  10
2001:DB8::76/128                                Local   Local     02d05h54m  0
       system                                                          0
2001:DB8:101::/48                               Remote  ISIS      00h04m25s  15
       FE80::E84B:FFFF:FE00:0-"toBNG"                                  20
2001:DB8:102::/48                               Remote  ISIS      00h04m25s  15
       FE80::E84B:FFFF:FE00:0-"toBNG"                                  20
2001:DB8:201::/48                               Remote  ISIS      00h04m25s  15
       FE80::E84B:FFFF:FE00:0-"toBNG"                                  20
2001:DB8:202::/48                               Remote  ISIS      00h04m25s  15
       FE80::E84B:FFFF:FE00:0-"toBNG"                                  20
2001:DEAD::/64                                  Local   Local     01h42m18s  0
       toBNG                                                           0
-------------------------------------------------------------------------------
```

```
No. of Routes: 7
Flags: L = LFA nexthop available    B = BGP backup route available
       n = Number of times nexthop is repeated
===============================================================================
A:R1#
```

Alternatively the same result can be achieved with OSPF/OSPFv3.

Option 2 – Export Policy (from protocol direct)

The key properties for the second option are:

- Subscriber interface subnets and prefixes are distributed into the network by applying an export policy.
- It is most typically used in combination with BGP based route distribution.
- It works with the numbered model only.

The following export policy is used for this example.

```
configure
    router
        policy-options
            policy-statement "local-subnets-out"
                entry 10
                    from
                        protocol direct
                    exit
                    action accept
                    exit
                exit
            exit
```

In this option the BNG relies on BGP using the policy local-subnets-out as an export policy. The neighbor address is the IPv4 system address of router R1.

```
configure
    router
        autonomous-system 65536
        bgp
            group "grp-1"
                family ipv4 ipv6
                export "local-subnets-out"
                peer-as 65536
                neighbor 192.0.2.76
                    advertise-label ipv6
                exit
            exit
            no shutdown
        exit
```

The following command shows the IPv4 routes advertised by applying the local-subnets-out policy. The subscriber interface subnets are advertised, as are some other local subnets.

```
*A:BNG# show router bgp neighbor 192.0.2.76 advertised-routes ipv4
===============================================================================
 BGP Router ID:192.0.2.75       AS:65536       Local AS:65536
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
===============================================================================
BGP IPv4 Routes
===============================================================================
Flag  Network                                              LocalPref   MED
      Nexthop                                              Path-Id     Label
      As-Path
-------------------------------------------------------------------------------
i     10.1.1.0/24                                          100         None
      192.0.2.75                                           None        -
      No As-Path
i     10.1.2.0/24                                          100         None
      192.0.2.75                                           None        -
      No As-Path
i     10.2.1.0/24                                          100         None
      192.0.2.75                                           None        -
      No As-Path
i     10.2.2.0/24                                          100         None
      192.0.2.75                                           None        -
      No As-Path
i     10.11.11.1/32                                        100         None
      192.0.2.75                                           None        -
      No As-Path
i     192.0.2.75/32                                        100         None
      192.0.2.75                                           None        -
      No As-Path
i     192.168.12.0/24                                      100         None
      192.0.2.75                                           None        -
      No As-Path
i     192.168.202.0/24                                     100         None
      192.0.2.75                                           None        -
      No As-Path
-------------------------------------------------------------------------------
Routes : 8
===============================================================================
*A:BNG#
```

The same applies for IPv6.

```
*A:BNG# show router bgp neighbor 192.0.2.76 advertised-routes ipv6
===============================================================================
 BGP Router ID:192.0.2.75       AS:65536       Local AS:65536
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
===============================================================================
```

```
BGP IPv6 Routes
===============================================================================
Flag  Network                                          LocalPref   MED
      Nexthop                                          Path-Id     Label
      As-Path
-------------------------------------------------------------------------------
i     2001:DB8::11/128                                 100         None
      ::FFFF:C000:24B                                  None        2
      No As-Path
i     2001:DB8::75/128                                 100         None
      ::FFFF:C000:24B                                  None        2
      No As-Path
i     2001:DB8:101::/48                                100         None
      ::FFFF:C000:24B                                  None        2
      No As-Path
i     2001:DB8:102::/48                                100         None
      ::FFFF:C000:24B                                  None        2
      No As-Path
i     2001:DB8:201::/48                                100         None
      ::FFFF:C000:24B                                  None        2
      No As-Path
i     2001:DB8:202::/48                                100         None
      ::FFFF:C000:24B                                  None        2
      No As-Path
i     2001:DEAD::/64                                   100         None
      ::FFFF:C000:24B                                  None        2
      No As-Path
-------------------------------------------------------------------------------
Routes : 7
===============================================================================
*A:BNG#
```

The corresponding IPv4 RIB on router R1 lists the subscriber-interface subnets, not the individual subscriber host addresses. Notice the list also includes other routes local to the BNG.

```
*A:R1# show router "Base" route-table ipv4
===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                             Type    Proto     Age         Pref
      Next Hop[Interface Name]                                     Metric
-------------------------------------------------------------------------------
10.1.1.0/24                                    Remote  BGP       00h13m34s   170
       192.168.12.1                                                0
10.1.2.0/24                                    Remote  BGP       00h13m34s   170
       192.168.12.1                                                0
10.2.1.0/24                                    Remote  BGP       00h13m34s   170
       192.168.12.1                                                0
10.2.2.0/24                                    Remote  BGP       00h13m34s   170
       192.168.12.1                                                0
10.11.11.1/32                                  Remote  BGP       00h13m34s   170
       192.168.12.1                                                0
192.0.2.75/32                                  Remote  ISIS      00h15m38s   15
       192.168.12.1                                                10
192.0.2.76/32                                  Local   Local     03h11m54s   0
       system                                                      0
```

```
192.168.12.0/24                                 Local   Local     03h11m25s  0
       toBNG                                                          0
192.168.202.0/24                                Remote  BGP       00h13m34s  170
       192.168.12.1                                                   0
-------------------------------------------------------------------------------
No. of Routes: 9
Flags: L = LFA nexthop available    B = BGP backup route available
       n = Number of times nexthop is repeated
===============================================================================
*A:R1#
```

The corresponding IPv6 RIB on router R1 lists the subscriber-interface prefixes, not the individual subscriber host addresses/prefixes. They are tunneled through the IPv4 core.

```
*A:R1# show router "Base" route-table ipv6
===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
2001:DB8::11/128                                Remote  BGP       00h00m49s  170
       192.0.2.75 (tunneled)                                          0
2001:DB8::75/128                                Remote  ISIS      00h54m05s  15
       FE80::E84B:FFFF:FE00:0-"toBNG"                                 10
2001:DB8::76/128                                Local   Local     05h18m12s  0
       system                                                         0
2001:DB8:101::/48                               Remote  BGP       00h00m49s  170
       192.0.2.75 (tunneled)                                          0
2001:DB8:102::/48                               Remote  BGP       00h00m49s  170
       192.0.2.75 (tunneled)                                          0
2001:DB8:201::/48                               Remote  BGP       00h00m49s  170
       192.0.2.75 (tunneled)                                          0
2001:DB8:202::/48                               Remote  BGP       00h00m49s  170
       192.0.2.75 (tunneled)                                          0
2001:DEAD::/64                                  Local   Local     05h17m42s  0
       toBNG                                                          0
-------------------------------------------------------------------------------
No. of Routes: 8
Flags: L = LFA nexthop available    B = BGP backup route available
       n = Number of times nexthop is repeated
===============================================================================
*A:R1# #
```

The same export policy can be used in combination with IGP based route distribution. However, when IGP based route distribution is needed option 1 is the preferred method.

Option 3 – Export Protocol (from protocol sub-mgmt)

The key properties for the third option are:

- Host addresses and prefixes are distributed into the network by applying an export policy.
- It is most typically used in combination with BGP based route distribution, as IGP based route distribution does not scale for a large number of subscribers.
- It is most typically used for the unnumbered model, and in some cases for the numbered model.

The following export policy is used for this option.

```
configure
    router
        policy-options
            policy-statement "subsc-hosts-out"
                entry 10
                    from
                        protocol sub-mgmt
                    exit
                    action accept
                    exit
                exit
            exit
```

In this option the BNG relies on BGP using the policy subsc-hosts-out as an export policy.

```
configure
    router
        autonomous-system 65536
        bgp
            group "grp-1"
                family ipv4 ipv6
                export "subsc-hosts-out"
                peer-as 65536
                neighbor 192.0.2.76
                    advertise-label ipv6
                exit
            exit
            no shutdown
        exit
```

The following command shows the IPv4 routes advertised by applying the subsc-hosts-out policy. Now the subscriber host addresses are advertised individually.

```
*A:BNG# show router bgp neighbor 192.0.2.76 advertised-routes ipv4
===============================================================================
 BGP Router ID:192.0.2.75       AS:65536       Local AS:65536
===============================================================================
 Legend -
```

```
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
===============================================================================
BGP IPv4 Routes
===============================================================================
Flag  Network                                              LocalPref   MED
      Nexthop                                              Path-Id     Label
      As-Path
-------------------------------------------------------------------------------
?     10.1.1.101/32                                        100         0
      192.0.2.75                                           None        -
      No As-Path
?     10.1.1.102/32                                        100         0
      192.0.2.75                                           None        -
      No As-Path
?     10.2.2.1/32                                          100         0
      192.0.2.75                                           None        -
      No As-Path
?     10.2.2.2/32                                          100         0
      192.0.2.75                                           None        -
      No As-Path
-------------------------------------------------------------------------------
Routes : 4
===============================================================================
*A:BNG#
```

For IPv6, the host addresses and prefixes are advertised.

```
*A:BNG# show router bgp neighbor 192.0.2.76 advertised-routes ipv6
===============================================================================
 BGP Router ID:192.0.2.75       AS:65536       Local AS:65536
===============================================================================
 Legend -
 Status codes  : u - used, s - suppressed, h - history, d - decayed, * - valid
 Origin codes  : i - IGP, e - EGP, ? - incomplete, > - best, b - backup
===============================================================================
BGP IPv6 Routes
===============================================================================
Flag  Network                                              LocalPref   MED
      Nexthop                                              Path-Id     Label
      As-Path
-------------------------------------------------------------------------------
?     2001:DB8:101::1/128                                  100         0
      ::FFFF:C000:24B                                      None        2
      No As-Path
?     2001:DB8:101:1::1/128                                100         0
      ::FFFF:C000:24B                                      None        2
      No As-Path
?     2001:DB8:102::/56                                    100         0
      ::FFFF:C000:24B                                      None        2
      No As-Path
?     2001:DB8:102:100::/56                                100         0
      ::FFFF:C000:24B                                      None        2
      No As-Path
-------------------------------------------------------------------------------
Routes : 4
===============================================================================
*A:BNG#
```

The corresponding IPv4 RIB on router R1 looks as follows. Notice the individual host addresses do appear.

```
A:R1# show router route-table ipv4
===============================================================================
Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age         Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.1.1.101/32                                   Remote  BGP       00h40m49s  170
       192.168.12.1                                                    0
10.1.1.102/32                                   Remote  BGP       00h19m53s  170
       192.168.12.1                                                    0
10.2.2.1/32                                     Remote  BGP       00h44m49s  170
       192.168.12.1                                                    0
10.2.2.2/32                                     Remote  BGP       00h44m17s  170
       192.168.12.1                                                    0
192.0.2.75/32                                   Remote  ISIS      00h59m41s  15
       192.168.12.1                                                    10
192.0.2.76/32                                   Local   Local     01h22m11s  0
       system                                                          0
192.168.12.0/24                                 Local   Local     01h21m42s  0
       toBNG                                                           0
-------------------------------------------------------------------------------
No. of Routes: 7
Flags: L = LFA nexthop available    B = BGP backup route available
       n = Number of times nexthop is repeated
===============================================================================
A:R1#
```

The corresponding IPv6 RIB on router R1 looks as follows. Notice the individual host addresses and prefixes are distributed in this case.

```
A:R1# show router route-table ipv6
===============================================================================
IPv6 Route Table (Router: Base)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age         Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
2001:DB8::76/128                                Local   Local     01h22m16s  0
       system                                                          0
2001:DB8:101::1/128                             Remote  BGP       00h36m40s  170
       192.0.2.75 (tunneled)                                           0
2001:DB8:101:1::1/128                           Remote  BGP       00h36m40s  170
       192.0.2.75 (tunneled)                                           0
2001:DB8:102::/56                               Remote  BGP       00h36m40s  170
       192.0.2.75 (tunneled)                                           0
2001:DB8:102:100::/56                           Remote  BGP       00h36m40s  170
       192.0.2.75 (tunneled)                                           0
2001:DEAD::/64                                  Local   Local     01h21m47s  0
       toBNG                                                           0
-------------------------------------------------------------------------------
No. of Routes: 6
Flags: L = LFA nexthop available    B = BGP backup route available
       n = Number of times nexthop is repeated
```

```
===============================================================================
A:R1#
```

Conclusion

This example explains how to configure and use the Routed CO model. The subscriber and the group interfaces were configured for the numbered, unnumbered and hybrid scenario, showing the flexibility of the Routed CO model in terms of subnet/prefix assignment as well as the impact on the forwarding and the reachability to and from the subscriber hosts.

Subscriber Redundancy for Routed CO

This chapter provides information about Subscriber Redundancy for Routed CO (SRRP).

Topics in this chapter include:

Applicability

This chapter is applicable to SROS routers and was initially based on release 7.0.R5. The CLI is updated to release 14.0.R4.

Summary

This chapter focuses on the delivery of redundant services in an enhanced subscriber management (ESM) routed-CO environment using Internet enhanced service (IES) or virtual private routed network (VPRN).

It is applicable to delivering high speed Internet (HSI), voice-over-IP (VoIP) and video-on-demand (VoD) to subscribers.

Redundancy is provided at two levels:

- system redundancy
- network redundancy

The system redundancy is based on the high availability features of the SR OS routers, such as component redundancy (power, fans, control processor modules etc.) and software redundancy (service and protocol redundancy and non-stop-routing), which are not discussed here.

The network redundancy for subscriber access in an ESM routed CO environment requires that each broadband service access node (BSAN) is dual-homed to two SR OS routers, either in a point-to-point fashion with the BSANs having direct physical connectivity to the SR OS routers, or by having Layer 2 aggregation between the BSANs and the SR OS routers.

This connection will operate in a master-slave relationship providing both link and system redundancy for the subscribers on the BSAN when accessing the configured services.

Subscriber redundancy for routed-CO aims to minimize the outage due to a failure.

This chapter provides configuration and troubleshooting commands for SRRP with **static-host ip-mac**.

Knowledge of the triple play service delivery architecture (TPSDA) concepts is assumed throughout this document.

Overview

There are three components on SR OS routers to implement network redundancy:

1. Redundant access from the BSANs to two SR OS routers provided by the SRRP.
2. Mirroring of subscriber state between two SR OS routers is achieved through the multi-chassis synchronization (MCS) protocol.
3. Backup spoke SDP traffic path between two SR OS peers.

These components are shown in Figure 219.

Figure 219 **Network Redundancy Components for ESM Routed CO**

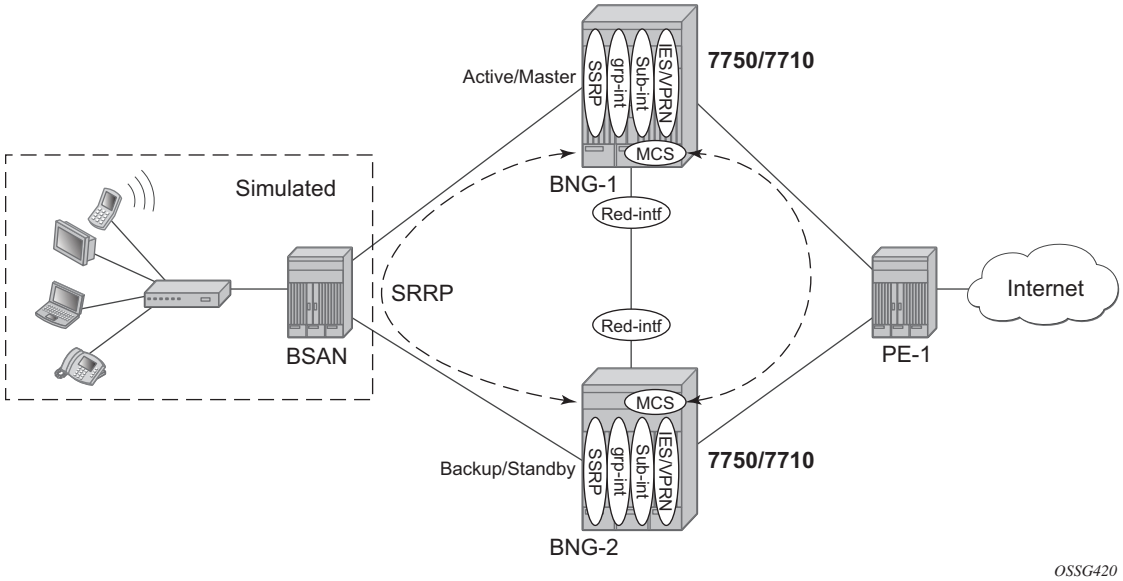

The following configuration tasks should be done first and are not explained in this chapter:

- Basic service router configuration (system interface, IGP, MPLS, BGP)
- Routed CO service topology: VPRN or IES service with subscriber and group interface on broadband network gateways (BNGs)
- ESM configuration
- Static host configuration

This chapter will focus on SRRP in a VPRN service subscriber-interface on BNG (routed-CO). In case of routed CO, it is also possible to configure SRRP in the base routing instance using an IES service.

SRRP Protocol

The SRRP protocol operates on a specific SAP within the group interface under the subscriber-interface of an IES/VPRN service. Through a method similar to the virtual router redundancy protocol (VRRP), it provides a set of default gateways to the subscribers on the BSAN. These are active on the SR OS router in the master state and inactive on the router in the backup state. Traffic is forwarded upstream and

downstream through the master. If the backup loses connectivity with the master (for example, fails to receive SRRP messages from the master), it transitions to the master state and takes over the ownership of the default gateways and the responsibility for forwarding traffic to and from the subscriber. This provides redundancy from the BSAN toward the provider network.

If an SRRP fail over were to occur, it is important that the subscriber state (IP/MAC addresses, QOS profiles, etc.) be immediately available on the new SRRP master; otherwise, subscriber traffic will be dropped due to the anti-spoofing security. The subscriber state is synchronized through the MCS protocol, which mirrors the subscriber state between the two peers. This allows both peers to know the details of the active subscribers and therefore forward traffic on their behalf with the correct QOS actions both to and from the BSAN if that peer is the SRRP master.

The last topic relates to the forwarding from the provider network to the subscribers. If the IP routing causes this traffic to forward through the SRRP master, then the traffic will automatically be forwarded to the subscriber. However, if the provider routing scheme causes traffic destined to a subscriber to arrive at the router in the SRRP backup state for that subscriber, it will be dropped as the backup does not forward traffic out of the subscriber SAPs.

To avoid this, a redundant interface is configured between the two SR OS routers under the subscriber/group-interfaces. Any traffic arriving on the router for an active subscriber, where its associated SRRP instance is in the backup state, will be forwarded through the redundant interface to the SRRP master, which in turn forwards the traffic to the subscriber.

In addition to successful forwarding the traffic, this operation also maintains the subscriber QOS compliance as all traffic for a given subscriber enters and exits the routed-CO interface through a single SAP, allowing the associated IOM hardware to perform the correct QOS actions.

Configuration

Subscriber Interface Configuration

Redundant Default Gateway

The redundant default gateway IP addresses must be configured under the subscriber-interface (within the IES/VPRN service) for each subnet defined.

Three subnets are configured under the subscriber-interface sub-int-1, each with a **gw-ip-address** which is uses as the default gateway by the subscribers in that subnet.

```
# on BNG-2
configure
    service
        vprn 1 customer 1 create
            --- snipped ---
            subscriber-interface "sub-int-1" create
                address 10.2.0.2/16 gw-ip-address 10.2.0.254
                address 10.3.0.2/16 gw-ip-address 10.3.0.254
                address 10.4.0.2/16 gw-ip-address 10.4.0.254
                --- snipped ---
            exit
        exit
    exit
exit
```

The **gw-ip-address** could be a virtual (unused) address in this subnet or the address of one of the actual subscriber-interfaces on the two routers. It must not be used as a subscriber address.

If DHCP were to be used, the associated subscriber-interface address should be used as the gi-address configured for DHCP under the group-interface (will not be covered here as static host is used). This ensures that the offer returned from the DHCP server and arriving at the SRRP backup (rather than master) will be forwarded by the backup SRRP router to the master SRRP router through the redundant interface.

In environments where there are many subscribers, it will take time to synchronize the subscriber state between the peers when the subscriber-interface is enabled (perhaps, after a reboot). In order to ensure that the state has time to get synchronized, a hold timer can be applied to the subscriber interface. The optional *init-only* parameter can be added to use this timer only after a reboot.

```
*A:BNG-2>config>service>vprn>red-if>hold-time#
 [no] down            - Configure the hold time when the interface is coming up
 [no] up              - Configure the hold time when the interface is going down

*A:BNG-2>config>service>vprn>red-if>hold-time#


# on BNG-2
configure
    service
        vprn 1 customer 1 create
            subscriber-interface "sub-int-1" create
                hold-time
                    down ip 1200 init-only
                exit
            exit
        exit
    exit
```

```
exit
```

Group Interface Configuration

The group interface **group-int-1** providing connectivity to the BSAN is configured under the subscriber interface:

```
# on BNG-2
configure
    service
        vprn 1 customer 1 create
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                --- snipped ---
            exit
        exit
    exit
exit
```

Static Host Configuration

Enable the sub-sla-mgmt and sub-ident-policy **sub-id-default** under **sap 1/1/3:1** and define static host (ip-mac) with **sla-profile sla-profile-1**, **sub-profile sla-profile-1** and **subscriber static-host-routed-10.2.0.3**.

```
# on BGG-2
configure
    service
        vprn 1 customer 1 create
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    sap 1/1/3:1 create
                        sub-sla-mgmt
                            sub-ident-policy "sub-id-default"
                            no shutdown
                        exit
                        static-host ip 10.2.0.3 mac 00:00:00:00:00:01 create
                            sla-profile "sla-profile-1"
                            sub-profile "sub-profile-1"
                            subscriber "static-host-routed-10.2.0.3"
                            no shutdown
                        exit
                    exit
                exit
            exit
        exit
    exit
exit
```

SRRP Configuration

In order for the redundant gateway information to be used by subscribers through SAPs belonging to a particular group-interface, an SRRP instance must be added in the group interface context.

```
# on BNG-2
configure
    service
        vprn 1 customer 1 create
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    srrp 1 create
                        --- snipped ---
                    exit
                exit
            exit
            no shutdown
        exit
    exit
exit
```

At this point, any subscriber ARPing for the gw-ip-address will receive a response from the SRRP master with a source MAC of 00-00-5E-00-01-<xx>, where <xx> is the first byte of the SRRP identifier in hexadecimal, so in this case for SRRP=1 the source MAC will be 00-00-5E-00-01-01.

The redundant default gateway MAC address could be explicitly configured, if desired, by use of the **gw-mac** parameter.

```
*A:BNG-2>config>service>vprn>sub-if>grp-if>srrp# gw-mac
  - gw-mac <mac-address>
  - no gw-mac
 <mac-address>        : xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx
*A:BNG-2>config>service>vprn>sub-if>grp-if>srrp#
```

There can be only one SRRP instance per group-interface and all SRRP identifiers must be unique per system.

When an SRRP instance is enabled, or when there is an SRRP failover from one device to another, gratuitous ARPs are sent on all VLANs associated with this instance for all gw-ip-addresses (for example, on all subscriber SAPs and on the SRRP message path SAP in the associated group interface). This allows all downstream devices to relearn the path to the new master.

The SRRP messages are normally forwarded through the BSAN, thereby verifying the connectivity from one router through the BSAN to the other router. In order to achieve this, a non-subscriber SAP must be configured under the **group-interface** which is referenced in the SRRP configuration by the **message-path** parameter. This not only selects the SAP to be used for the SRRP messages but also avoids the subscriber anti-spoofing from automatically dropping the received messages (as there would be no subscriber IP-to-MAC entry corresponding to the received information) and causing both peers to become master.

The message-path SAP configuration effectively disables IP-MAC anti-spoofing on that SAP.

```
# on BNG-2
configure
    service
        vprn 1 customer 1 create
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    ---- snipped ---
                    sap 1/1/3:2 create
                    exit
                    srrp 1 create
                        message-path 1/1/3:2
                        no shutdown
                    exit
                exit
            exit
            no shutdown
        exit
    exit
exit
```

The SRRP messages are then not sent within the same SAP as the subscriber data traffic, but it is assumed that the path traversed by the SRRP messages is the same as would be used for the subscriber data; if this is not the case, then the SRRP state would not necessarily reflect a failure in the data path.

The master of the SRRP instance generates advertisement messages at the keep-alive-interval (which is encoded in the message) ranging from one (1) to 100 in multiples of 100ms with a default of 10 (for example, 1 second). The SRRP backup will monitor the reception of these messages and assume the role of the master if three consecutive messages are not received.

At all times the keep-alive-interval of the master is used.

```
*A:BNG-2>config>service>vprn>sub-if>grp-if>srrp# keep-alive-interval
  - keep-alive-interval <interval>
  - no keep-alive-interval

 <interval>           : [1..100] tenths of a second

*A:BNG-2>config>service>vprn>sub-if>grp-if>srrp#
```

Only two devices can participate in an SRRP protocol exchange for a given SRRP instance, this being another difference from VRRP which allows more potential backup devices. This is a consequence of the direct relationship between the SRRP instance and the associated redundant interface and MCS peering.

This protocol exchange is also used for the master/backup election, based on the priority (1 to 254) configured in the SRRP instance. The device with the highest priority will become master.

```
*A:BNG-2>config>service>vprn>sub-if>grp-if>srrp# priority
  - no priority
  - priority <priority>

 <priority>           : [1..254]

*A:BNG-2>config>service>vprn>sub-if>grp-if>srrp#
```

The message source IP address (system IP address) is used as a tie break when the priorities are the same (the lower IP address becomes the master). The master/backup status is per SRRP instance (not per IP address). For example, the master is the active gateway for all gw-ip-addresses under the subscriber interface for the associated group-interface (this is true even if the backup is the IP address owner for one of the gw-ip-addresses). Priority 0 is sent by the master when it is transitioning to the backup role due to the appearance of a high priority peer. Higher priority backups always preempt a lower priority master.

A basic form of load distribution can be achieved by having the master SRRP for some group-interfaces on one peer and the master for other group interfaces on the other peer. Clearly, a failure may cause all masters to be active on a single peer, which must be taken into account when designing the network.

The minimum keep-alive-timer of 1 is configured, together with the message-path giving the SAP to be used for the SRRP messages. The priority is set to 250 (default is 100) in order to force this peer to be the SRRP master when both peers are active.

```
# on BNG-2
configure
    service
        vprn 1 customer 1 create
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                --- snipped ---
                srrp 1 create
                    keep-alive-interval 1
                    message-path 1/1/3:2
                    priority 250
                    no shutdown
                exit
            exit
        exit
    exit
exit
```

SRRP Configuration Notes

A VRRP policy statement can be added to the SRRP instance definition in order to dynamically adjust the SRRP priority based on certain non-SRRP related events occurring (for example, port down, LAG degrade, host unreachable or route unknown).

The gw-ip-addresses are accessible by active subscribers, for example, regardless which peer is the master, an active subscriber can ping or telnet to its associated gw-ip-address (clearly, filters can be configured to control this).

Bi-directional Forwarding Detection

Bi-directional forwarding detection (BFD) can be configured with SRRP to speed up the convergence.

```
# on BNG-2
configure
    service
        vprn 1 customer 1 create
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                --- snipped ---
                    srrp 1 create
                        --- snipped ---
                        bfd-enable 2 interface "bfd-1" dst-ip 10.1.1.1
                        no shutdown
                    exit
                exit
            exit
        exit
    exit
exit
```

An IES service needs to be created for the BFD session.

```
# on BNG-1
configure
    service
        ies 2 customer 1 create
            interface "bfd-1" create
                address 10.1.1.1/31
                bfd 100 receive 100 multiplier 3
                sap 1/1/3:3 create
                exit
            exit
            no shutdown
        exit
    exit
exit
```

Monitoring In-Band Communications Path

In order to monitor the in-band communications path between the subscribers and two routers, SRRP uses a slightly modified VRRP advertisement message.

The SRRP message destination IP address (224.0.0.18) and IP protocol number (112) are the same as for VRRP but there are changes in the following areas:

- The source IP address of the message is the system IP address, as opposed to the interface IP address.
- The protocol version has been set to eight (8) (the current version of VRRP is two (2)).
- The virtual router identifier has been extended from one byte (maximum 255) to four bytes (maximum 4294967295) though the maximum number of SRRPs that can be defined is 255.
- The source MAC address is included instead of the virtual router IP addresses (this being 00-00-5E-00-01-<xx>, where <xx> is the first byte of the SRRP identifier in hexadecimal).

Synchronizing the SRRP Peer State

In order to troubleshoot an SRRP environment, the state of each peer is synchronized with the other peer through multi-chassis synchronization (MCS). MCS is a proprietary protocol used for synchronizing application state between two peers. SRRP will function without synchronizing its state but this synchronization allows for the current state of both the local and remote peers to be displayed and appropriate error messages to be reported when the peer state is not correct. It also allows the master SRRP subscriber-interface to be pinged through the backup peer (through the redundant interface).

To link information being mirrored between two routers, a **sync-tag** value is configured to correspond to either an entire port/LAG or under a port/LAG for a VLAN range. This allows each router to know exactly which information should be in sync on each device. The sync-tag must be unique on the two peers involved.

This example configuration shows only the SRRP aspects. The SRRP instance has been configured for MCS peer 192.0.2.1 using VLANs 1-2 on port 1/1/3. Here, a sync-tag is associated with the SRRP instance.

```
# on BNG-2
configure
    redundancy
        multi-chassis
```

```
            peer 192.0.2.1 create
                authentication-key "sync-testing"
                sync
                    srrp
                    sub-mgmt ipoe
                    port 1/1/3 "st1" create
                        range 1-2 sync-tag "st1"
                    exit
                    no shutdown
                exit
                no shutdown
            exit
        exit
    exit
exit
```

Alternatively, if the information needs to be synchronized for all VLANs on port 1/1/3, then the following port command could be used instead of the preceding port-plus-ranges shown.

```
# on BNG-2
configure
    redundancy
        multi-chassis
            peer 192.0.2.2 create
                authentication-key "sync-testing"
                sync
                    srrp
                    sub-mgmt ipoe
                    port 1/1/3 sync-tag "st1" create
                    exit
                    no shutdown
                exit
                no shutdown
            exit
        exit
    exit
exit
```

The VLANs used within the group interfaces must match between the two peers, clearly the physical ports identifiers may differ.

Multi-Chassis Synchronization

Multi-chassis synchronization (MCS) is a general propriety protocol used to synchronize information between two peers. It can be used for the several applications, such as:

- IGMP
- IGMP snooping

- Subscriber management
- Subscriber router redundancy protocol

This chapter only covers the subscriber management and SRRP applications.

Subscriber Management Synchronization

In order to ensure that the QOS defined for a subscriber is adhered to, all traffic for a given subscriber needs to be forwarded by a single port. When an MSAN is dual-homed to two routers, this is achieved using the SRRP protocol (described above) and the redundant interface (described below); specifically, the traffic is forwarded through the SRRP master of the related group-interface.

When a subscriber is created on the master SRRP, a host route (/32) for its IP address is inserted in the FIB pointing towards the appropriate group-interface.

The same subscriber on the backup peer will have a host route in the IP FIB pointing to the redundant interface. On the backup peer, this requires the subscriber subnet to also be present in the FIB, which in turn requires one of the following:

- The local subscriber-interface is up.
- The subscriber subnets are learned from the active broadband network gateway (BNG) through the routing protocol.
- Forcing the subscriber-interface to stay up by creating a dummy group interface with the **oper-up-when-empty** command.

```
# on BNG-2
configure
    service
        vprn 1 customer 1 create
            subscriber-interface "sub-int-1" create
                group-interface "group-int-1" create
                    oper-up-while-empty
                    sap 1/1/3:1 create
                        --- snipped ---
                    exit
                exit
            exit
        exit
    exit
exit
```

Redundant Interface

The requirement for the redundant interface comes from the situation where traffic destined to a subscriber arrives on the router but the associated SRRP state is not master.

When the SRRP state is backup for a particular group-interface, subscriber traffic is normally not forwarded in/out of the associated subscriber SAPs. Also traffic could arrive but the specific group-interface for that subscriber is down. These situations could occur due to the regular routing within the provider network or temporarily during an SRRP failover. Note that as the subscriber subnets are configured under the subscriber-interface, it is not possible to stop advertising these subnets into the provider core in the case where only a subset of the group interfaces are down or the associated SRRP instances are in backup. The advertisement of the subscriber subnet could therefore attract traffic to the router, while not being the SRRP master.

In these cases, traffic must be sent to the active SRRP router in order to be forwarded to the subscriber. This is achieved through the configuration of a redundant interface between two SRRP peers, protecting against failures of related group interfaces.

The redundant interface is configured under the IES/VRPN service. It must use a single pseudowire, configured as a spoke SDP, to provide connectivity to the peer router. This is essential as it avoids any possibility of the traffic being misrouted by any other system between the two peers. It can either use a /31 IP subnet mask or a longer mask with the remote IP being explicitly specified.

```
# on BRG-2
configure
    service
        vprn 1 customer 1 create
            --- snipped ---
            redundant-interface "bng-2-bng-1-vprn-1" create
                address 192.168.4.1/31
                spoke-sdp 21:1 create
                exit
            exit
            --- snipped ---
        exit
    exit
exit
```

If a non /31 address is used, the remote IP address will be required.

```
*A:BNG-2>config>service>vprn# redundant-interface "bng-2-bng-1-vprn-1"
*A:BNG-2>config>service>vprn>red-if# address 192.168.4.1/30
INFO: PIP #1399 Invalid or unspecified remote IP address - Non /
31 address requires remote IP address
```

The remote-ip address can be defined by the following command:

```
*A:BNG-2>config>service>vprn>red-if# address 192.168.4.1/30 remote-ip 192.4.1.2
```

Each group-interface must then be associated with a redundant interface.

```
# on BNG-2
configure
    service
        vprn 1 customer 1 create
            --- snipped ---
            subscriber-interface "sub-int-1" create
                --- snipped ---
                group-interface "group-int-1" create
                    redundant-interface "bng-2-bng-1-vprn-1"
                    oper-up-while-empty
                    --- snipped ---
                exit
            exit
            no shutdown
        exit
    exit
exit
```

Only one redundant interface is required for a given IES/VPRN service on each peer, though it is possible to create multiple redundant interfaces and assign group interfaces to each individually. Clearly, each redundant interface needs to terminate on a matching redundant interface in the corresponding peer service.

When traffic arrives from the core network for an active subscriber on a SAP in group-interface **group-int-1**, and if the associated SRRP instance is in the backup state, then this traffic will be forwarded over the **redundant-interface bng-2-bng-1-vprn-1** to the peer router. It will then be forwarded to the subscriber as the associated SRRP instance will be in the master state.

The information about the redundant interfaces is mirrored through MCS as part of the SRRP synchronization.

Show and Debug Commands

Routing Table Related Information

A host route (/32) for the static host is inserted in the FIB pointing toward the appropriate group-interface.

```
*A:BNG-2# show router 1 route-table protocol sub-mgmt

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age         Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.2.0.1/32                                     Remote  Sub Mgmt  06d01h29m  0
       [bng-2-bng-1-vprn-1]                                          0
10.2.0.3/32                                     Remote  Sub Mgmt  06d01h49m  0
       [group-int-1]                                                 0
10.3.0.1/32                                     Remote  Sub Mgmt  06d01h29m  0
       [bng-2-bng-1-vprn-1]                                          0
10.4.0.1/32                                     Remote  Sub Mgmt  06d01h29m  0
       [bng-2-bng-1-vprn-1]                                          0
-------------------------------------------------------------------------------
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
*A:BNG-2#
```

The same subscriber on the backup peer will have a host route in the FIB pointing to the redundant interface.

```
*A:BNG-1# show router 1 route-table protocol sub-mgmt

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age         Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.2.0.2/32                                     Remote  Sub Mgmt  02h44m34s  0
       [bng-1-bng-2-vprn-1]                                          0
10.2.0.3/32                                     Remote  Sub Mgmt  04h10m39s  0
       [bng-1-bng-2-vprn-1]                                          0
10.3.0.2/32                                     Remote  Sub Mgmt  02h44m34s  0
       [bng-1-bng-2-vprn-1]                                          0
10.4.0.2/32                                     Remote  Sub Mgmt  02h44m34s  0
       [bng-1-bng-2-vprn-1]                                          0
-------------------------------------------------------------------------------
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
*A:BNG-1#
```

Subscriber Related Information

To check the subscriber information:

```
*A:BNG-2# show service id 1 subscriber-hosts

===============================================================
Subscriber Host table
===============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID Origin       Fwding State
---------------------------------------------------------------
1/1/3:1                static-host-routed-10.2.0.3
  10.2.0.3
    00:00:00:00:00:01    N/A       Static       Fwding
---------------------------------------------------------------
Number of subscriber hosts : 1
===============================================================
*A:BNG-2#
```

To check the subscriber details:

```
*A:BNG-2# show service id 1 subscriber-hosts mac 00:00:00:00:00:01 detail

===============================================================
Subscriber Host table
===============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID Origin       Fwding State
---------------------------------------------------------------
1/1/3:1                static-host-routed-10.2.0.3
  10.2.0.3
    00:00:00:00:00:01    N/A       Static       Fwding
---------------------------------------------------------------
Subscriber-interface  : sub-int-1
Group-interface       : group-int-1
Sub Profile           : sub-profile-1
SLA Profile           : sla-profile-1
App Profile           : N/A
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : EA4DFF0000000057AB1A60
Acct-Q-Inst-Session-Id: EA4DFF0000000157AB1A60
Address Origin        : LUDB
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
DIAMETER session ID Gx: N/A
---------------------------------------------------------------
Number of subscriber hosts : 1
===============================================================
*A:BNG-2#
```

The same command on the peer would show the same information except for the port identifier part of the SAP, which is specific to the peer and may differ.

MCS Redundancy Related Information

The high-level state of MCS can be seen in the following output:

```
*A:BNG-2# show redundancy multi-chassis all

===============================================================================
Multi-Chassis Peers
===============================================================================
Peer Info                    Client    Admin        Oper         State
-------------------------------------------------------------------------------
Peer Address : 192.0.2.1
Source Addre*: 192.0.2.2
Peer Admin   : inService
Authenticati*: None
* indicates that the corresponding row element may have been truncated.
                             MC-Sync:  inService    inService    inSync
                             MC-Ring:  --           --           --
                             MC-Endpt: --           --           --
                             MC-Lag:   outOfService outOfService --
                             MC-IPsec: --           --           Disabled
===============================================================================
*A:BNG-2#
```

Information about the peering, the use of authentication, the state of MCS (Enabled) and the fact that MCS is inSync is shown.

```
*A:BNG-2# show redundancy multi-chassis sync

===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 192.0.2.1
Description             : (Not Specified)
Authentication          : Disabled
Source IP Address       : 192.0.2.2
Admin State             : Enabled
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     : SUBMGMT-IPOE SUBMGMT-PPPOE SRRP
Sync Admin State        : Up
Sync Oper State         : Up
Sync Oper Flags         :
DB Sync State           : inSync
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
===============================================================================
===============================================================================
*A:BNG-2#
```

In the output, it can be seen that SUBMGMT-IPOE, SUBMGMT-PPPOE and SRRP are client applications and they are inSync with 26 database entries both on this peer and the remote peer.

If the preceding command requested detailed output for peer 192.0.2.1, additional information would be shown.

```
*A:BNG-2# show redundancy multi-chassis sync peer 192.0.2.1 detail

===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 192.0.2.1
Description             : (Not Specified)
Authentication          : Disabled
Source IP Address       : 192.0.2.2
Admin State             : Enabled
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     : SUBMGMT-IPOE SUBMGMT-PPPOE SRRP
Sync Admin State        : Up
Sync Oper State         : Up
Sync Oper Flags         :
DB Sync State           : inSync
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
===============================================================================
MCS Application Stats
===============================================================================
Application             : igmp
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
```

```
--- snipped ---

-------------------------------------------------------------------------------
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
-------------------------------------------------------------------------------

--- snipped ---

===============================================================================
Ports synced on peer 192.0.2.1
===============================================================================
Port/Encap                    Tag
-------------------------------------------------------------------------------
1/1/3                         st1
===============================================================================

--- snipped ---

===============================================================================
Diameter proxy instances synced on peer 192.0.2.1
===============================================================================
Diameter-Peer-Policy             Tag
-------------------------------------------------------------------------------
No instances found
===============================================================================
===============================================================================
*A:BNG-2#
```

This shows that there are 26 entries on both peers for SRRP.

If the hold-time parameter is configured under the subscriber-interface, in order to allow time for the MCS to fully synchronize, its setting and current expiry time can be seen as follows:

```
A:BNG-2# show service id 1 interface "sub-int-1" detail

===============================================================================
Interface Table
===============================================================================

-------------------------------------------------------------------------------
Interface
-------------------------------------------------------------------------------
If Name          : sub-int-1
Admin State      : Up                   Oper (v4/v6)      : Down/--
```

```
Down Reason Code : delayedStartEnabled
Down Reason V4   : delayedStartEnabled
Down Reason V6   : ifProtoOperDown
Protocols        : None
IP Addr/mask     : 10.2.0.2/16
HoldUp-Time      : 0                    Track Srrp Inst   : 0
IP Addr/mask     : 10.3.0.2/16
HoldUp-Time      : 0                    Track Srrp Inst   : 0
IP Addr/mask     : 10.4.0.2/16
HoldUp-Time      : 0                    Track Srrp Inst   : 0
Ignore Port State: None
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
Description      : (Not Specified)
If Index         : 5                    Virt. If Index    : 5
Last Oper Chg    : 08/18/2016 10:22:20  Global If Index   : 258
Mon Oper Grp     : None
Srrp En Rtng     : Disabled             Hold time         : N/A
V4 Delay IfUp    : 300 init-only        V4 Time to IfUp   : 278
                                        Unmatching Subnet : No
                                        Unmatching Pfxs   : No
If Type          : VPRN Sub

DHCP Details
Gi-Addr          : Not configured       Gi-Addr as Src Ip : Disabled
Virt. subnet     : disabled

===============================================================================
Interface sub-int-1 group-interfaces
===============================================================================
Interface-Name                   Adm       Opr(v4/v6)  Mode    Port/SapId
   IP-Address                                                   PfxState
-------------------------------------------------------------------------------
group-int-1                      Up        Up/--       VPRN G* 1/1/3
-------------------------------------------------------------------------------
Group-Interfaces : 1
===============================================================================
* indicates that the corresponding row element may have been truncated.
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
A:BNG-2#
```

Tool Dump Commands Related Information

The database entries can be view in more detail with the **tools dump redundancy multi-chassis** command.

```
*A:BNG-2# tools dump redundancy multi-chassis sync-database

The following totals are for:
 peer ip ALL, port/lag/sdp ALL, sync-tag ALL, application ALL
Valid Entries:                  26
Locally Deleted Entries:        0
```

```
Locally Deleted Alarmed Entries: 0
Pending Global Delete Entries:   0
Omcr Alarmed Entries:            0
Omcr Standby Entries:            0
Associated Shared Records (ALL): 0
Associated Shared Records (LD):  0
*A:BNG-2#
```

The output of the “tool dump redundancy multi-chassis” command is as follows:

```
*A:BNG-2# tools dump redundancy multi-chassis srrp-sync-database
Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
   SMK_BASE_CONFIG / 192.0.2.1
Data Info:
    srrpId: 1    svcId: 1     svcType: VPRN
    system IP: 0xc0000201    Group interface MAC: ea:4c:01:01:00:03
    Gateway MAC: 00:00:5e:00:01:01
    Subscriber interface name: sub-int-1

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
   SMK_BASE_CONFIG / 192.0.2.2
Data Info:
    srrpId: 1    svcId: 1     svcType: VPRN
    system IP: 0xc0000202    Group interface MAC: ea:4d:01:01:00:03
    Gateway MAC: 00:00:5e:00:01:01
    Subscriber interface name: sub-int-1

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
   SMK_GRP_IF / 192.0.2.1
Data Info:
    Group interface name: group-int-1
    Redundant interface name: bng-1-bng-2-vprn-1
    Redundant Interface IP/Mask:
        IP: 0xc0a80400 Mask: 0xfffffffe
    AdminUp: Up, OperState: SRRP_STATE_BACKUP_SHUNT, InUsePriority: 100, Red-If OK:
 Yes, MessageSap OK: Yes

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
   SMK_GRP_IF / 192.0.2.2
Data Info:
    Group interface name: group-int-1
    Redundant interface name: bng-2-bng-1-vprn-1
    Redundant Interface IP/Mask:
        IP: 0xc0a80401 Mask: 0xfffffffe
    AdminUp: Up, OperState: SRRP_STATE_MASTER, InUsePriority: 250, Red-If OK: Yes,
 MessageSap OK: Yes

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
   SMK_GRP_IF_SAP_BUCKET0 / 192.0.2.1
Data Info:

--- snipped ---
```

```
Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
   SMK_GRP_IF_SAP_BUCKET9 / 192.0.2.2
Data Info:

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
   SMK_SUBNET_INFO / 192.0.2.1 vRtrId 2, ifIdx 5
Data Info:
    Subscriber IP Addr: 10.2.0.1     Mask: 0xffff0000     Gateway: 10.2.0.254
    Subscriber IP Addr: 10.3.0.1     Mask: 0xffff0000     Gateway: 10.3.0.254
    Subscriber IP Addr: 10.4.0.1     Mask: 0xffff0000     Gateway: 10.4.0.254

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
   SMK_SUBNET_INFO / 192.0.2.2 vRtrId 2, ifIdx 5
Data Info:
    Subscriber IP Addr: 10.2.0.2     Mask: 0xffff0000     Gateway: 10.2.0.254
    Subscriber IP Addr: 10.3.0.2     Mask: 0xffff0000     Gateway: 10.3.0.254
    Subscriber IP Addr: 10.4.0.2     Mask: 0xffff0000     Gateway: 10.4.0.254

*A:BNG-2#
```

Also shown are the port/VLANs synchronized with their respective sync-tags.

To further troubleshoot and debug this configuration, there are commands to both dump the sync and SRRP MCS information and to dump the SRRP database:

```
tools dump redundancy multi-chassis sync-database [application {sub-mgmt|srrp}]
```

The command provides the same information as the equivalent show commands. However, the detailed version gives more information about the contents of the sync-database.

For SRRP, there are entries for the base configuration, group interface and subnet information for each of the SRRP instances. This should show corresponding entries for the local and remote peer. Specifying the **sync-tag st1** shows only the information for SRRP instance 1.

```
*A:BNG-2# tools dump redundancy multi-chassis sync-database application srrp sync-
tag st1 detail

If no entries are present for an application, no detail will be displayed.

FLAGS LEGEND: ld - local delete; da - delete alarm; pd - pending global delete;
              oal - omcr alarmed; ost - omcr standby

Peer Ip 192.0.2.1


Application SRRP
Sap-id             Client Key
 SyncTag                          DLen  Flags           timeStamp
  deleteReason code and description                  #ShRec
-------------------------------------------------------------------------------
```

```
1/1/3:2               SMK_BASE_CONFIG / 192.0.2.1
 st1                              88    -- -- -- --- --- 08/10/2016 14:23:22
  0x0                                                   0
1/1/3:2               SMK_BASE_CONFIG / 192.0.2.2
 st1                              88    -- -- -- --- --- 08/10/2016 14:23:21
  0x0                                                   0
1/1/3:2               SMK_GRP_IF / 192.0.2.1
 st1                              212   -- -- -- --- --- 08/16/2016 15:57:11
  0x0                                                   0
1/1/3:2               SMK_GRP_IF / 192.0.2.2
 st1                              212   -- -- -- --- --- 08/16/2016 15:57:11
  0x0                                                   0
1/1/3:2               SMK_GRP_IF_SAP_BUCKET0 / 192.0.2.1
 st1                              4     -- -- -- --- --- 08/10/2016 14:23:22
  0x0                                                   0

  --- snipped ---

1/1/3:2               SMK_GRP_IF_SAP_BUCKET9 / 192.0.2.2
 st1                              4     -- -- -- --- --- 08/10/2016 14:23:21
  0x0                                                   0
1/1/3:2               SMK_SUBNET_INFO / 192.0.2.1 vRtrId 2, ifIdx 5
 st1                              40    -- -- -- --- --- 08/10/2016 14:23:22
  0x0                                                   0
1/1/3:2               SMK_SUBNET_INFO / 192.0.2.2 vRtrId 2, ifIdx 5
 st1                              40    -- -- -- --- --- 08/10/2016 14:23:21
  0x0                                                   0
The following totals are for:
 peer ip ALL, port/lag/sdp ALL, sync-tag st1, application SRRP
Valid Entries:                   26
Locally Deleted Entries:         0
Locally Deleted Alarmed Entries: 0
Pending Global Delete Entries:   0
Omcr Alarmed Entries:            0
Omcr Standby Entries:            0
Associated Shared Records (ALL): 0
Associated Shared Records (LD):  0
*A:BNG-2#
```

This same information can be seen in more detail by dumping the SRRP database. This output is for SRRP instance 1, and shows the detailed information for each peer. This should clearly reflect the configuration and current state of the SRRP instances. Again there are two entries (one for the local peer and the other for the remote peer) for the BASE_CONFIG, GRP_IF and SUBNET_INFO.

```
*A:BNG-2# tools dump redundancy multi-chassis srrp-sync-database instance 1
Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
   SMK_BASE_CONFIG / 192.0.2.1
Data Info:
    srrpId: 1    svcId: 1     svcType: VPRN
    system IP: 0xc0000201    Group interface MAC: ea:4c:01:01:00:03
    Gateway MAC: 00:00:5e:00:01:01
    Subscriber interface name: sub-int-1

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
```

```
   SMK_BASE_CONFIG / 192.0.2.2
Data Info:
    srrpId: 1    svcId: 1     svcType: VPRN
    system IP: 0xc0000202    Group interface MAC: ea:4d:01:01:00:03
    Gateway MAC: 00:00:5e:00:01:01
    Subscriber interface name: sub-int-1

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
   SMK_GRP_IF / 192.0.2.1
Data Info:
    Group interface name: group-int-1
    Redundant interface name: bng-1-bng-2-vprn-1
    Redundant Interface IP/Mask:
        IP: 0xc0a80400 Mask: 0xfffffffe
    AdminUp: Up, OperState: SRRP_STATE_BACKUP_SHUNT, InUsePriority: 100, Red-If OK:
 Yes, MessageSap OK: Yes

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
   SMK_GRP_IF / 192.0.2.2
Data Info:
    Group interface name: group-int-1
    Redundant interface name: bng-2-bng-1-vprn-1
    Redundant Interface IP/Mask:
        IP: 0xc0a80401 Mask: 0xfffffffe
    AdminUp: Up, OperState: SRRP_STATE_MASTER, InUsePriority: 250, Red-If OK: Yes,
 MessageSap OK: Yes

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
   SMK_GRP_IF_SAP_BUCKET0 / 192.0.2.1
Data Info:

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
   SMK_GRP_IF_SAP_BUCKET0 / 192.0.2.2
Data Info:

--- snipped ---

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
   SMK_GRP_IF_SAP_BUCKET9 / 192.0.2.1
Data Info:

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
   SMK_GRP_IF_SAP_BUCKET9 / 192.0.2.2
Data Info:

Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
   SMK_SUBNET_INFO / 192.0.2.1 vRtrId 2, ifIdx 5
Data Info:
    Subscriber IP Addr: 10.2.0.1     Mask: 0xffff0000     Gateway: 10.2.0.254
    Subscriber IP Addr: 10.3.0.1     Mask: 0xffff0000     Gateway: 10.3.0.254
    Subscriber IP Addr: 10.4.0.1     Mask: 0xffff0000     Gateway: 10.4.0.254
```

```
Tag Key:  sap = 1/1/3:2
Key Info: (Type/Owner)
   SMK_SUBNET_INFO / 192.0.2.2 vRtrId 2, ifIdx 5
Data Info:
    Subscriber IP Addr: 10.2.0.2     Mask: 0xffff0000     Gateway: 10.2.0.254
    Subscriber IP Addr: 10.3.0.2     Mask: 0xffff0000     Gateway: 10.3.0.254
    Subscriber IP Addr: 10.4.0.2     Mask: 0xffff0000     Gateway: 10.4.0.254

*A:BNG-2#
```

The following is an example of messages that could be seen due to this synchronization which otherwise would not be available.

An event will be generated in log 99, if the IP address was removed from the redundant interface on the remote peer.

```
*A:BNG-2# show log log-id 99

===============================================================================
Event Log 99
===============================================================================
Description : Default System Log
Memory Log contents  [size=500   next event=4  (not wrapped)]

3 2016/08/16 16:07:17.94 CEST WARNING: MC_REDUNDANCY #2012 vprn1 SRRP/MCS: Peer Red
 i/f down
"SRRP ID 1: Redundant interface bng-1-bng-2-vprn-1 on peer 192.0.2.1 /
 interface group
-int-1 does not match local 192.0.2.2 / interface group-int-1."

2 2016/08/16 16:07:17.94 CEST WARNING: MC_REDUNDANCY #2012 vprn1 SRRP/MCS: Peer Red
 i/f no addr
"SRRP ID 1: Redundant interface bng-1-bng-2-vprn-1 on peer 192.0.2.1 /
 interface group
-int-1 does not match local 192.0.2.2 / interface group-int-1."

1 2016/08/16 15:59:05.52 CEST INDETERMINATE: LOGGER #2010 Base Clear LOGGER
"Clear function clearLogId has been run with parameters: log-
id="99" context="".  The
 completion result is: success.  Additional error text, if any, is: "
*A:BNG-2#
```

The SRRP Instance Related Information

The SRRP instance information can be displayed by the following commands:

The master BNS shows the **master** in the operation status.

```
*A:BNG-2# show srrp

===============================================================================
SRRP Table
===============================================================================
```

```
ID        Service        Group Interface                Admin     Oper
-------------------------------------------------------------------------------
1         1              group-int-1                    Up        master
-------------------------------------------------------------------------------
No. of SRRP Entries: 1
===============================================================================
*A:BNG-2#
```

The backup BNG shows a **backupShunt** in the operation status.

```
*A:BNG-1# show srrp

===============================================================================
SRRP Table
===============================================================================
ID        Service        Group Interface                Admin     Oper
-------------------------------------------------------------------------------
1         1              group-int-1                    Up        backupShunt
-------------------------------------------------------------------------------
No. of SRRP Entries: 1
===============================================================================
*A:BNG-1#
```

To check detailed information:

```
*A:BNG-2# show srrp 1 detail

===============================================================================
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           : 192.0.2.2
Service ID          : VPRN 1
Group If            : group-int-1        MAC Address      : ea:4d:01:01:00:03
Grp If Description  : N/A
Grp If Admin State  : Up                 Grp If Oper State: Up
Subscriber If       : sub-int-1
Sub If Admin State  : Up                 Sub If Oper State: Up
Address             : 10.2.0.2/16        Gateway IP       : 10.2.0.254
Address             : 10.3.0.2/16        Gateway IP       : 10.3.0.254
Address             : 10.4.0.2/16        Gateway IP       : 10.4.0.254
Redundant If        : bng-2-bng-1-vprn-1
Red If Admin State  : Up                 Red If Oper State: Up
Address             : 192.168.4.1/31
Red Spoke-sdp       : 21:1
Msg Path SAP        : 1/1/3:2
Admin Gateway MAC   :                    Oper Gateway MAC : 00:00:5e:00:01:01
Config Priority     : 250                In-use Priority  : 250
Master Priority     : 250
Keep-alive Interval : 1 deci-seconds     Master Since     : 08/16/2016 15:57:11
Fib Population Mode : all
VRRP Policy 1       : None               VRRP Policy 2    : None

-------------------------------------------------------------------------------
```

```
BFD interface
-------------------------------------------------------------------------------
Service ID          : 2
Interface Name      : bfd-1
Src IP              : 10.1.1.0
Dst IP              : 10.1.1.1
Session Oper State  : connected

-------------------------------------------------------------------------------
Statistics
-------------------------------------------------------------------------------
Become Master       : 6                 Master Changes   : 11
Become Bkup Routing : 0                 Become Bkup Shunt: 6
Become Non-Master   : 5
Adv Sent            : 1882496           Adv Received     : 3358878
Pri 0 Pkts Sent     : 5                 Pri 0 Pkts Rcvd  : 0
Preempt Events      : 0                 Preempted Events : 5
Mesg Intvl Discards : 0                 Mesg Intvl Errors: 0

===============================================================================
*A:BNG-2#
```

If this command is executed on the backup, an extra line appears after the keep-alive-interval showing the interval during which the receipt of no SRRP messages would cause the master to be considered down, together with the instantaneous time to this interval expiring.

```
*A:BNG-1# show srrp 1 detail

===============================================================================
SRRP Instance 1
===============================================================================
Description         : (Not Specified)
Admin State         : Up                Oper State       : backupShunt
Preempt             : yes               One GARP per SAP : no
Monitor Oper Group  : None
System IP           : 192.0.2.1
Service ID          : VPRN 1
Group If            : group-int-1       MAC Address      : ea:4c:01:01:00:03
Grp If Description  : N/A
Grp If Admin State  : Up                Grp If Oper State: Up
Subscriber If       : sub-int-1
Sub If Admin State  : Up                Sub If Oper State: Up
Address             : 10.2.0.1/16       Gateway IP       : 10.2.0.254
Address             : 10.3.0.1/16       Gateway IP       : 10.3.0.254
Address             : 10.4.0.1/16       Gateway IP       : 10.4.0.254
Redundant If        : bng-1-bng-2-vprn-1
Red If Admin State  : Up                Red If Oper State: Up
Address             : 192.168.4.0/31
Red Spoke-sdp       : 12:1
Msg Path SAP        : 1/1/3:2
Admin Gateway MAC   :                   Oper Gateway MAC : 00:00:5e:00:01:01
Config Priority     : 100               In-use Priority  : 100
Master Priority     : 250
Keep-alive Interval : 1 deci-seconds    Master Since     : 08/16/2016 15:57:12
Master Down Interval: 0.300 sec (Expires in 0.250 sec)
Fib Population Mode : all
```

```
VRRP Policy 1       : None              VRRP Policy 2    : None

-------------------------------------------------------------------------------
BFD interface
-------------------------------------------------------------------------------
Service ID          : 2
Interface Name      : bfd-1
Src IP              : 10.1.1.1
Dst IP              : 10.1.1.0
Session Oper State  : connected

-------------------------------------------------------------------------------
Statistics
-------------------------------------------------------------------------------
Become Master       : 10                Master Changes   : 20
Become Bkup Routing : 4                 Become Bkup Shunt: 14
Become Non-Master   : 10
Adv Sent            : 3361406           Adv Received     : 1879862
Pri 0 Pkts Sent     : 6                 Pri 0 Pkts Rcvd  : 0
Preempt Events      : 5                 Preempted Events : 10
Mesg Intvl Discards : 0                 Mesg Intvl Errors: 0

===============================================================================
*A:BNG-1#
```

Monitoring the Traffic on Redundant Interface

The Oper State reflects both the state of the SRRP instance and its action with respect to the redundant interface. Specifically, when the peer is SRRP master the operational state is always master – traffic is sent directly to the subscriber over its associated SAP. If the peer is SRRP backup and the redundant interface is Up then the Oper State will be backupShunt, if the redundant interface is down then the Oper State is **backupRouting**. In the **backupShunt** state, traffic to the subscriber is shunted (for example, forwarded) across the redundant interface to the peer (to the master) in order to be forwarded to the subscriber.

When in the **backupRouting** state, the SRRP instance is in backup but the redundant interface is down, so the traffic is forwarded directly to the subscriber through its associated SAP.

A useful command to see the traffic on the redundant interface is:

```
*A:BNG-2# monitor service id 1 sdp 21:1 rate interval 11 repeat 3

===============================================================================
Monitor statistics for Service 1 SDP binding 21:1
===============================================================================
-------------------------------------------------------------------------------
At time t = 0 sec (Base Statistics)
-------------------------------------------------------------------------------
I. Fwd. Pkts.   : 1029                  I. Dro. Pkts.   : 0
E. Fwd. Pkts.   : 21                    E. Fwd. Octets  : 1830
```

```
-------------------------------------------------------------------------------
At time t = 11 sec (Mode: Rate)
-------------------------------------------------------------------------------
I. Fwd. Pkts.   : 1                    I. Dro. Pkts.   : 0
E. Fwd. Pkts.   : 0                    E. Fwd. Octets  : 0

-------------------------------------------------------------------------------
At time t = 22 sec (Mode: Rate)
-------------------------------------------------------------------------------
I. Fwd. Pkts.   : 2                    I. Dro. Pkts.   : 0
E. Fwd. Pkts.   : 0                    E. Fwd. Octets  : 0

-------------------------------------------------------------------------------
At time t = 33 sec (Mode: Rate)
-------------------------------------------------------------------------------
I. Fwd. Pkts.   : 1                    I. Dro. Pkts.   : 0
E. Fwd. Pkts.   : 0                    E. Fwd. Octets  : 0

===============================================================================
*A:BNG-2#
```

BFD-Related Information

To check the BFD session state.

```
*A:BNG-2# show router bfd session

===============================================================================
Legend:  wp = Working path   pp = Protecting path
===============================================================================
BFD Session
===============================================================================
If/Lsp Name/Svc-Id/RSVP-sess  State                Tx Intvl  Rx Intvl  Multipl
  Rem Addr/Info/SdpId:VcId    Protocols            Tx Pkts   Rx Pkts   Type
   LAG port                      LAG ID
-------------------------------------------------------------------------------
bfd-1                         Up                   100       100       3
   10.1.1.1                   srrp                 103787    103599    iom
-------------------------------------------------------------------------------
No. of BFD sessions: 1
===============================================================================
*A:BNG-2#
```

To check the MAC addresses of the SRRP, this is can be done by checking the MACs table in the BSAN.

```
*A:DSLAM# show service fdb-mac

===============================================================================
Service Forwarding Database
===============================================================================
ServId    MAC               Source-Identifier        Type     Last Change
                                                     Age
```

```
-------------------------------------------------------------------------------
1         00:00:00:00:00:01 sap:1/1/4:1              L/0      08/16/16 14:34:55
1         00:00:5e:00:01:01 sap:1/1/2:1              L/0      08/16/16 15:57:10
2         00:00:5e:00:01:01 sap:1/1/2:2              L/0      08/16/16 15:57:10
3         ea:4c:01:01:00:03 sap:1/1/1:3              L/0      08/16/16 14:34:53
3         ea:4d:01:01:00:03 sap:1/1/2:3              L/0      08/16/16 14:34:53
-------------------------------------------------------------------------------
No. of Entries: 5
-------------------------------------------------------------------------------
Legend:  L=Learned O=Oam P=Protected-MAC C=Conditional S=Static
===============================================================================
*A:DSLAM#
```

The first entry shows the source MAC address of SRRP 1 which is learned through the SRRP messages received from the SRRP master for that instance. The next two entries relate to the subscriber traffic to and from sub1. The router source MAC address is that of the SRRP instance.

The last two entries are the source MAC address BFD received from each BNG.

SRRP Debug Commands

There are debug command to show the SRRP protocol events and packets.

```
*A:BNG-2# debug router 1 srrp events
*A:BNG-2# debug router 1 srrp packets
```

To display the debugging information, a dedicated log should be created:

```
# on BNG-2
debug
    router "1"
        srrp
            packets
            events
        exit
    exit
exit
```

The following output displays a sample SRRP debug log:

```
1917 2016/08/16 16:13:16.09 CEST MINOR: DEBUG #2001 vprn1 SRRP
"SRRP: Sending Pkt

Version (SRRP)     : 8
Type               : Advertisement (1)
Vr Id              : 1
Priority           : 250
Count Ip Addresses : 3
Advertise Interval : 10 centi-second
Checksum           : 0x25ef
```

```
Raw Pkt:

81 00 fa 00 00 00 00 01 00 00 5e 00 01 01 00 0a
00 03 25 ef "
```

As an example, in the following output BNG-2 is the SRRP master for the first instance (1) and sends SRRP advertisement messages. Then BNG-2 receives an SRRP message with a higher priority (254) from its peer BNG-1. This causes an event **Become Pending-Backup Shunt** where BNG-2 prepares to transition to the backup state. To achieve this, BNG-2 sends and SRRP message with priority 0. If BNG-2 continues to receive SRRP messages (with priority 254) from its peer BNG-1, it passes into the backup state with the event **Become Backup Shunt**.

```
2062 2016/08/16 16:13:30.59 CEST MINOR: DEBUG #2001 vprn1 SRRP
"SRRP: Sending Pkt

Version (SRRP)     : 8
Type               : Advertisement (1)
Vr Id              : 1
Priority           : 250
Count Ip Addresses : 3
Advertise Interval : 10 centi-second
Checksum           : 0x25ef

Raw Pkt:

81 00 fa 00 00 00 00 01 00 00 5e 00 01 01 00 0a
00 03 25 ef "


2063 2016/08/16 16:13:30.59 CEST MINOR: DEBUG #2001 vprn1 SRRP
"SRRP: Receiving Pkt

Version (SRRP)     : 8
Type               : Advertisement (1)
Vr Id              : 1
Priority           : 254
Count Ip Addresses : 3
Advertise Interval : 10 centi-second
Checksum           : 0x21ee

Raw Pkt:

81 01 fe 00 00 00 00 01 00 00 5e 00 01 01 00 0a
00 03 21 ee "


2064 2016/08/16 16:13:30.59 CEST MINOR: DEBUG #2001 vprn1 SRRP
"SRRP: Event
Become Pending-Backup Shunt: vRtrId 2, ifIdx 6, IPv4 vr_id 1, Master IP 192.0.2.
1"


2065 2016/08/16 16:13:30.59 CEST MINOR: DEBUG #2001 vprn1 SRRP
"SRRP: Sending Pkt
```

```
Version (SRRP)     : 8
Type               : Advertisement (1)
Vr Id              : 1
Priority           : 0
Count Ip Addresses : 3
Advertise Interval : 10 centi-second
Checksum           : 0x1ff0

Raw Pkt:

81 00 00 00 00 00 00 01 00 00 5e 00 01 01 00 0a
00 03 1f f0 "


2066 2016/08/16 16:13:30.59 CEST MINOR: DEBUG #2001 vprn1 SRRP
"SRRP: Receiving Pkt

Version (SRRP)     : 8
Type               : Advertisement (1)
Vr Id              : 1
Priority           : 254
Count Ip Addresses : 3
Advertise Interval : 10 centi-second
Checksum           : 0x21ef

Raw Pkt:

81 00 fe 00 00 00 00 01 00 00 5e 00 01 01 00 0a
00 03 21 ef "


2067 2016/08/16 16:13:30.59 CEST MINOR: DEBUG #2001 vprn1 SRRP
"SRRP: Event
Become Backup Shunt: vRtrId 2, ifIdx 6, IPv4 vr_id 1, Master IP 192.0.2.1"


2068 2016/08/16 16:13:30.59 CEST MINOR: DEBUG #2001 vprn1 SRRP
"SRRP: Receiving Pkt

Version (SRRP)     : 8
Type               : Advertisement (1)
Vr Id              : 1
Priority           : 254
Count Ip Addresses : 3
Advertise Interval : 10 centi-second
Checksum           : 0x21ef

Raw Pkt:

81 00 fe 00 00 00 00 01 00 00 5e 00 01 01 00 0a
00 03 21 ef "
```

The output for the SRRP message as captured with tshark.

```
Frame 2 (60 bytes on wire, 60 bytes captured)
    Arrival Time: Aug 16, 2016 15:24:52.362531000
    [Time delta from previous captured frame: 0.019443000 seconds]
```

```
    [Time delta from previous displayed frame: 0.000000000 seconds]
    [Time since reference or first frame: 0.019443000 seconds]
    Frame Number: 2
    Frame Length: 60 bytes
    Capture Length: 60 bytes
    [Frame is marked: False]
    [Protocols in frame: eth:vlan:ip:vrrp]
Ethernet II, Src: IETF-VRRP-virtual-router-VRID_01 (00:00:5e:00:01:01), Dst:
 IPv4mcast_00:00:12 (01:00:5e:00:00:12)
    Destination: IPv4mcast_00:00:12 (01:00:5e:00:00:12)
        Address: IPv4mcast_00:00:12 (01:00:5e:00:00:12)
        .... ...1 .... .... .... .... = IG bit: Group address (multicast/broadcast)
        .... ..0. .... .... .... .... = LG bit: Globally unique address (factory def
ault)
    Source: IETF-VRRP-virtual-router-VRID_01 (00:00:5e:00:01:01)
        Address: IETF-VRRP-virtual-router-VRID_01 (00:00:5e:00:01:01)
        .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
        .... ..0. .... .... .... .... = LG bit: Globally unique address (factory def
ault)
    Type: 802.1Q Virtual LAN (0x8100)
802.1Q Virtual LAN, PRI: 0, CFI: 0, ID: 2
    000. .... .... .... = Priority: 0
    ...0 .... .... .... = CFI: 0
    .... 0000 0000 0010 = ID: 2
    Type: IP (0x0800)
    Trailer: 0000
Internet Protocol, Src: 192.0.2.2 (192.0.2.2), Dst: 224.0.0.18 (224.0.0.18)
    Version: 4
    Header length: 20 bytes
    Differentiated Services Field: 0xc0 (DSCP 0x30: Class Selector 6; ECN: 0x00)
        1100 00.. = Differentiated Services Codepoint: Class Selector 6 (0x30)
        .... ..0. = ECN-Capable Transport (ECT): 0
        .... ...0 = ECN-CE: 0
    Total Length: 40
    Identification: 0x167e (5758)
    Flags: 0x00
        0... = Reserved bit: Not set
        .0.. = Don't fragment: Not set
        ..0. = More fragments: Not set
    Fragment offset: 0
    Time to live: 255
    Protocol: VRRP (0x70)
    Header checksum: 0x0213 [correct]
        [Good: True]
        [Bad : False]
    Source: 192.0.2.2 (192.0.2.2)
    Destination: 224.0.0.18 (224.0.0.18)
Virtual Router Redundancy Protocol
    Version 8, Packet type 1 (Advertisement)
        1000 .... = VRRP protocol version: 8
        .... 0001 = VRRP packet type: Advertisement (1)
    Virtual Rtr ID: 0
    Priority: 251 (Non-default backup priority)
    Count IP Addrs: 0
    Auth Type: No Authentication (0)
    Adver Int: 0
    Checksum: 0x0001 [correct]
```

SRRP Traffic Marking

The SRRP messages are sent by default with DSCP of **nc1** and with 802.1p bits of **0**, as can be seen in the following tshark snippet.

```
802.1Q Virtual LAN, PRI: 0, CFI: 0, ID: 2
    000. .... .... .... = Priority: 0
    ...0 .... .... .... = CFI: 0
    .... 0000 0000 0010 = ID: 2
    Type: IP (0x0800)
    Trailer: 0000
Internet Protocol, Src: 192.0.2.2 (192.0.2.2), Dst: 224.0.0.18 (224.0.0.18)
    Version: 4
    Header length: 20 bytes
    Differentiated Services Field: 0xc0 (DSCP 0x30: Class Selector 6; ECN: 0x00)
        1100 00.. = Differentiated Services Codepoint: Class Selector 6 (0x30)
        .... ..0. = ECN-Capable Transport (ECT): 0
        .... ...0 = ECN-CE: 0
    Total Length: 40


*A:BNG-2# show qos dscp-table

============================================================
DSCP Mapping
============================================================
DSCP Name      DSCP Value     TOS (bin)      TOS (hex)
------------------------------------------------------------
be             0              0000 0000      00
--- snipped ---
nc1            48             1100 0000      C0
--- snipped ---
cp63           63             1111 1100      FC
============================================================
*A:BNG-2#
```

Where DSCP 0x30=48 (DSCP value). This can be changed to EF, for example, using the following command:

```
*A:BNG-2# configure service vprn 1 sgt-qos application srrp dscp ef
```

Conclusion

This chapter provides configuration and troubleshooting commands for SRRP with static (IP-MAC) host in a Layer 3 Routed-CO (IES/VPRN subscriber interface) context.

Virtual Residential Gateway Authentication Scenarios

This chapter describes virtual residential gateway authentication scenarios.

Topics in this chapter include:

Applicability

This chapter is applicable to SR OS routers and is based on SR OS 14.0.R3.

Overview

In the virtual residential gateway (vRGW) model, the Layer 3 (L3) functions are moved out of the traditional residential gateway (RGW) and into the network. The Virtual Residential Gateway Home Pool Management chapter provides the rationale for this scenario, and describes how services must be configured for the service router to support this model.

The home network can be self-managed through a service portal, where end users can connect and change home-specific settings; see Figure 220. The portal logic is implemented in a controlling entity providing a RESTful interface. A protocol conversion and mediation platform (PCMP) is needed to translate the RESTful interface into RADIUS (and SNMP), and vice versa. The PCMP operates in conjunction with the 5620 SAM. In the remainder of this chapter, PCMP and the controller are represented as a single component.

Managing individual BRGs requires the vRGW to maintain a context for every BRG, so each BRG is identified through the BRG ID. This context is created when authenticating the BRG, and stores the home-level settings. The BRG context is deleted when the BRG is not deemed alive anymore.

The BRG ID can be derived from any of the parameters in the RADIUS Access-Request message, such as the SAP, tunnel-source, or called-station ID by a controlling entity; see the Implicit Authentication and Explicit Authentication sections that follow in this chapter. Typically, the MAC address of the BRG serves as the BRG ID, and that is what is used throughout this chapter.

Figure 220 **BRG and Home Device Management**

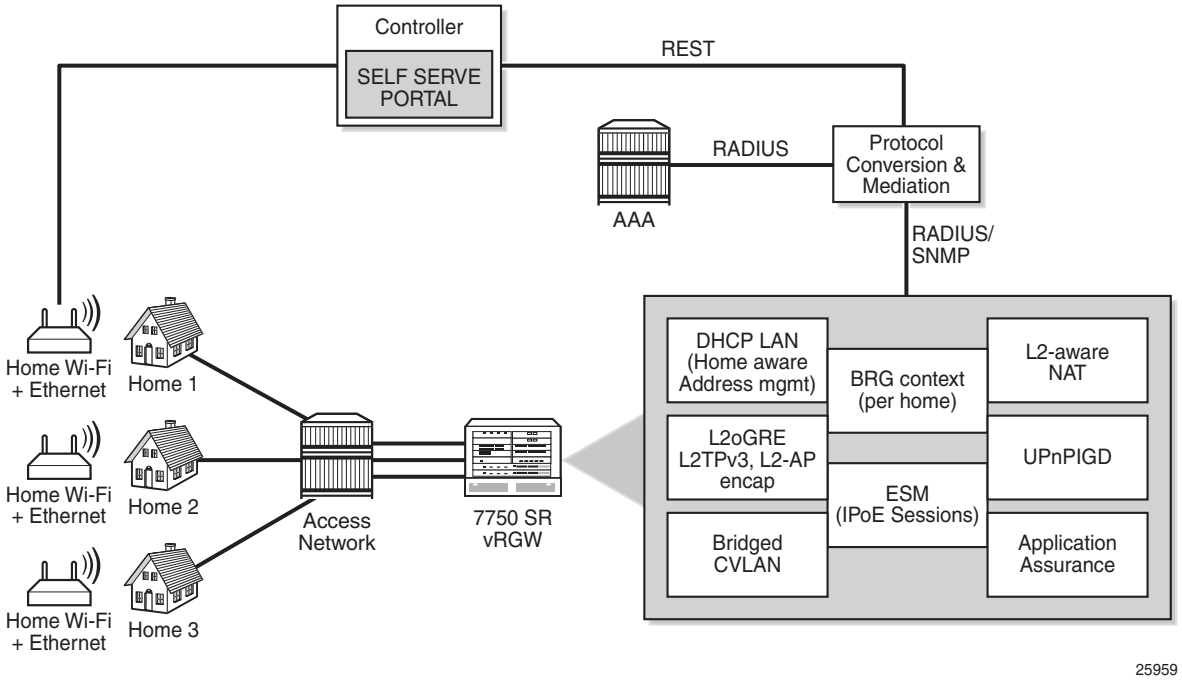

The vRGW supports two scenarios for authenticating bridged gateways and their hosts:

- Explicit authentication
- Implicit authentication

Explicit Authentication

Two main phases are distinguished in the explicit authentication scenario; see Figure 221. The first phase is the BRG authentication phase, the second phase is the device authentication phase. The example in Figure 221 uses IPv4, but also works with IPv6.

Figure 221 **Explicit BRG Authentication**

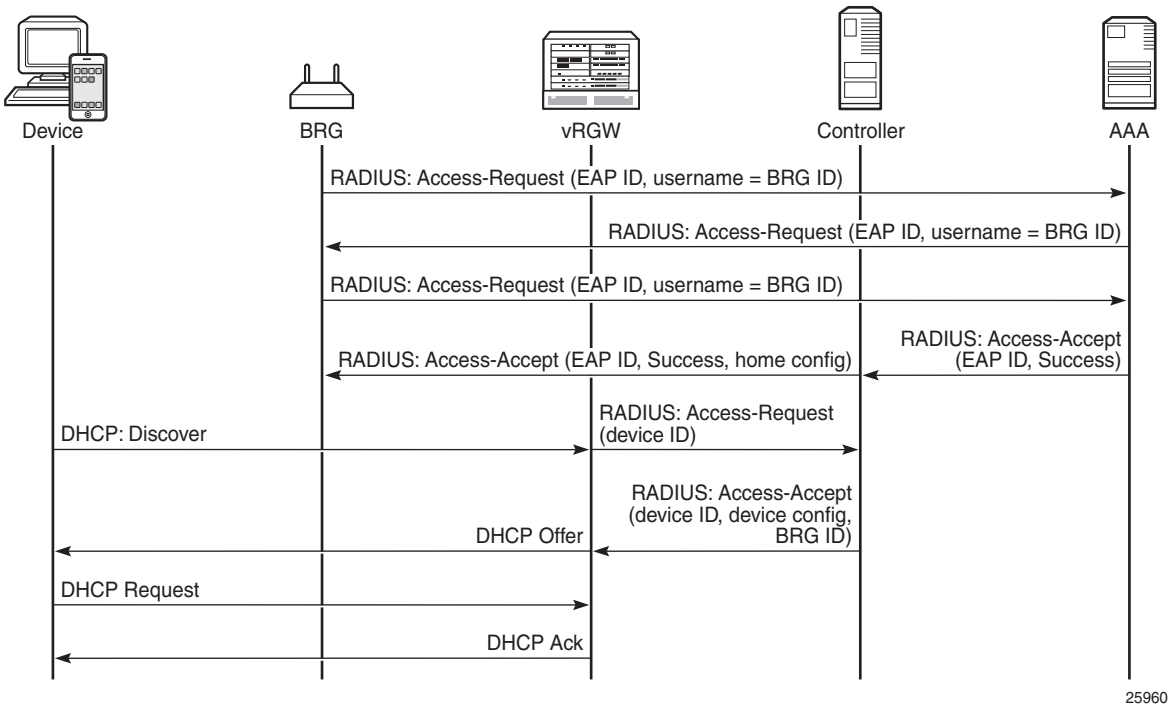

The first message in the first phase is an Access-Request message sent by the BRG toward the AAA/RADIUS server, and uses the extensible authentication protocol (EAP). This message is proxied by the vRGW as well as by the controller to the AAA server, and the BRG ID is used as the username. The last message of the first phase is the Access-Accept message. When the controller receives this message from the AAA server, it fetches and adds the per-home configuration parameters to the Access-Accept message before forwarding this message to the vRGW.

The second phase starts with the Discover message of a typical Discover-Offer-Request-Ack (DORA) message sequence. The vRGW then initiates device authentication toward the controller, which returns the BRG ID and, optionally, device-specific configuration data in an Access-Accept message. The controller usually will not proxy device authentication toward the AAA server. Typically, the RADIUS protocol is used between the vRGW and the controller. As such, the vRGW will send an Access-Request message to the controller for every new device, including static devices, to get the per-device configuration.

Usually, the vRGW combines the home-specific data with the device-specific data, where the more specific device data overrules the home-specific data. The combined data is then used to create the corresponding ESM hosts and IPoE sessions.

If the home-specific data includes Alc-Reserved-Addresses attributes defining static devices, these devices are authenticated automatically as soon as at least one dynamic device is connected.

The explicit authentication scenario requires the BRG profile to include a RADIUS proxy server; see the Configuration section in this chapter for a practical example.

Implicit Authentication

In the implicit authentication scenario, the BRG is authenticated when the first device connects to the vRGW; see Figure 222 for an example.

Figure 222 **Implicit BRG Authentication**

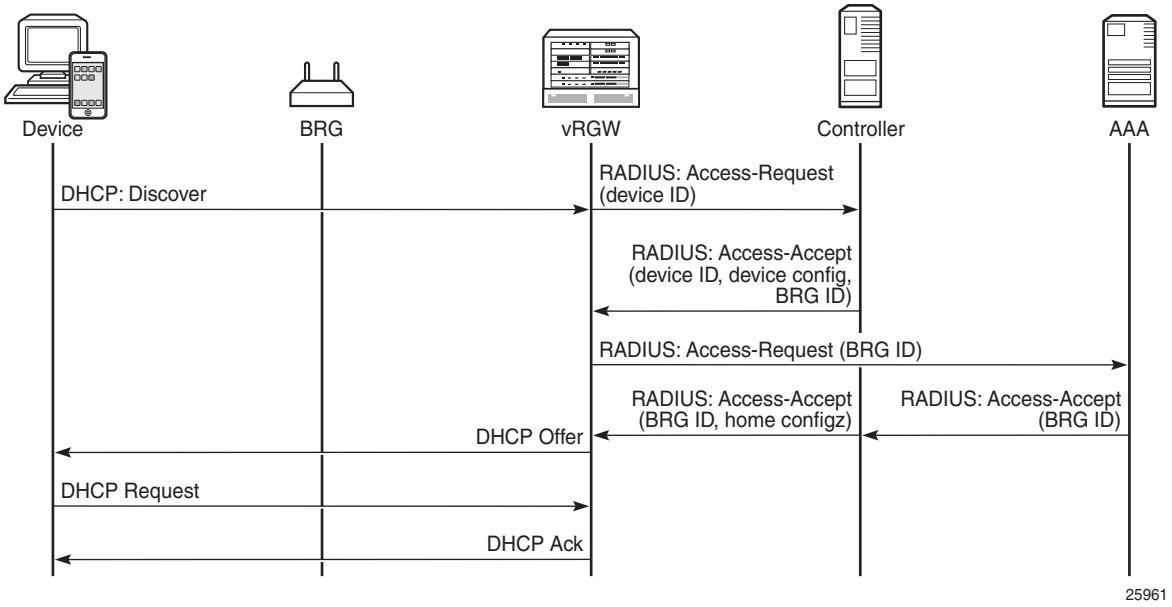

In Figure 222, the Discover message triggers the vRGW to send an Access-Request message toward the controller. The controller returns device-specific data including the BRG ID. Because there is no context for this BRG yet, the vRGW starts BRG authentication toward the AAA/RADIUS server. The controller proxies this message, and on return adds the home-specific data to the Access-Accept message. The overall result is that the vRGW now knows the home- and device-specific data.

As with the explicit authentication scenario, the combined data is then used to create the corresponding ESM hosts and IPoE sessions.

No separate BRG authentication is required when subsequent devices connect and device authentication returns a BRG ID that is already known to the vRGW.

As with the explicit authentication scenario, if the home-specific data includes Alc-Reserved-Addresses attributes defining static home devices, the static devices are authenticated automatically when at least one dynamic device is connected.

The implicit authentication scenario requires the BRG profile to include a RADIUS authentication context defining a RADIUS server policy and a password; see the Configuration section in this chapter for an example.

Connectivity Verification and BRG Deletion

For the purpose of clearing resources when these resources are not needed anymore, the vRGW performs the BRG connectivity verification and deletion process. When BRG connectivity is considered lost, the BRG and its hosts are deleted automatically.

Parameters controlling the BRG connectivity verification and deletion process are located in the BRG profile context:

```
*A:BNG>config>subscr-mgmt>brg-profile#
   connectivity-verification count 3 timeout 30 retry-time 900
     count <nr-of-attempts>     : [1..5]       - default: 3
     timeout <timeout-seconds>  : [5..60]      - default: 30
     retry-time <retry-seconds> : [300..3600]  - default: 300
   hold-time <seconds> [30..86400]             - default: no hold-time
   initial-hold-time <seconds> [0..900]        - default: 300
```

When the last dynamic host associated with a BRG is deleted, while at the same time connectivity-verification is enabled, the vRGW starts a liveliness test toward the BRG through ICMP (v4 or v6) messages. These messages are sent to either the BRG tunnel source IP address or the BRG RADIUS source IP address. If the BRG has neither of these addresses (for example, because each BRG is managed through a unique VLAN in the implicit authentication scenario), connectivity verification is not executed and only the hold-timer applies. If no answer is returned by the BRG in time (timeout), the vRGW considers the BRG in a failed state, and tries again (retry-time) for some maximum number of times (count).

The vRGW starts the hold-timer when the maximum number of tries is reached, or when connectivity verification is disabled (no connectivity verification). When the hold-timer expires, the BRG context is deleted together with the associated hosts.

The initial-hold-timer is required in the scenario where operators want to use explicit authentication without connectivity-verification, and with no hold-timer defined. In this scenario, defining a non-zero initial-hold-timer value avoids BRG contexts from being deleted immediately after their creation.

If a new dynamic host connects while executing the liveliness test or while the hold-timer or the initial-hold-timer is running, the connectivity verification and BRG deletion process is canceled.

The hold-timer is ignored when manually clearing BRGs and related hosts with the following command.

```
clear subscriber-mgmt brg gateway
    brg-id <brg-ident>
        host <ieee-address>
        all-hosts
        idle-bindings
    all-gateways
```

Configuration

The services configuration shown in Figure 4 applies to the examples throughout this chapter. Because the main focus of this chapter is on BRG and home device authentication, the detailed configuration of VPLS 10 containing the capture SAP and of VPRN-2 containing the outside L2-aware NAT range is not repeated here. See the Virtual Residential Gateway Home Pool Management chapter for those configurations.

In summary, VPRN-1 provides the connection toward the customer premises, and hosts the NAT inside addresses. VPRN-2 provides the connection toward the Internet, and hosts the NAT outside addresses. Also, VPRN-3 provides connection to the management interfaces of the BRGs, and is used for the explicit BRG authentication scenario.

Figure 223 **Example Service Configuration for Explicit and Implicit BRG Authentication**

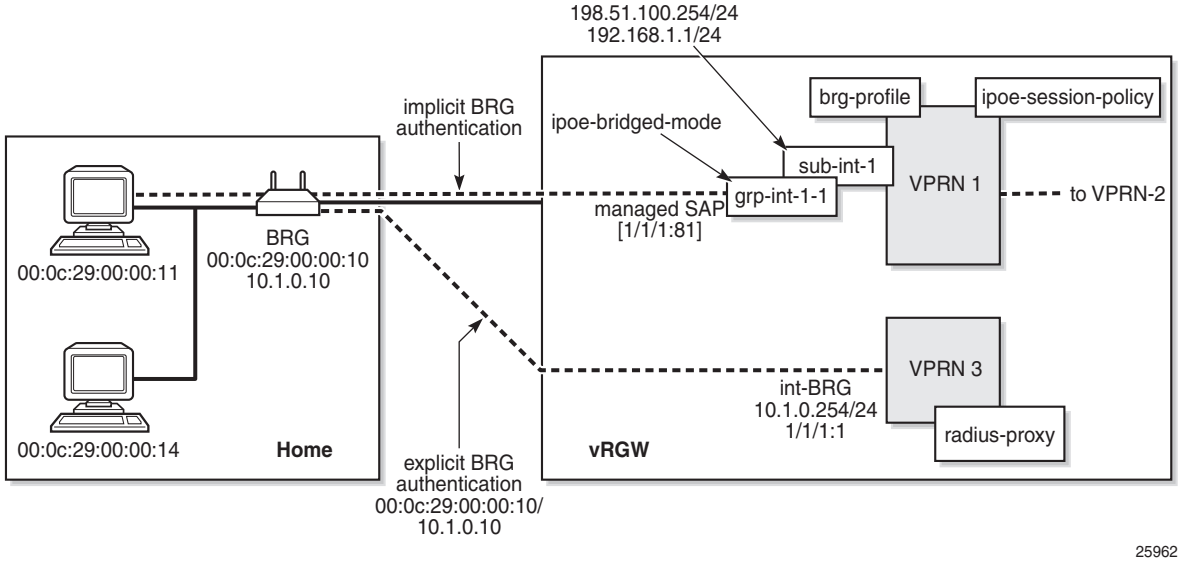

Service Configuration

An excerpt of the configuration of VPRN-1 follows. The group interface *grp-int-1-1* has RADIUS authentication enabled through authentication policy *radius-AUTH*, and BRG authentication through the default BRG profile *brg-prof-1*.

```
configure
    service
        vprn 1 customer 1 create
            --- snipped ---
            subscriber-interface "sub-int-1" create
                address 198.51.100.254/24
                address 192.168.1.1/24
                --- snipped ---
                group-interface "grp-int-1-1" create
                    ipv6
                        --- snipped ---
                        ipoe-bridged-mode
                    exit
                    --- snipped ---
                    authentication-policy "radius-AUTH"
                    ipoe-session
                        ipoe-session-policy "sess-pol-SAP-MAC"
                        sap-session-limit 128
                        no shutdown
                    exit
                    brg
                        default-brg-profile "brg-prof-1"
                        no shutdown
```

```
                    exit
                    oper-up-while-empty
                exit
            exit
            nat
                inside
                    l2-aware
                        address 192.168.0.1/16
                    exit
                exit
            exit
            no shutdown
        exit
    exit
exit
```

BRG Profile

The BRG profile *brg-prof-1* is defined in the subscriber management context and provides an SLA profile, a subscriber profile, a DHCP pool, a RADIUS server policy plus the corresponding password, and a RADIUS proxy server.

For explicit BRG authentication, the RADIUS proxy server is used; for implicit BRG authentication, the RADIUS server policy and password defined in the RADIUS authentication context are used.

```
configure
    subscriber-mgmt
        --- snipped ---
        brg-profile "brg-prof-1" create
            description "default BRG-profile, demo purposes"
            sla-profile-string "sla-prof-1"
            sub-profile-string "sub-prof-1"
            dhcp-pool
                subnet 192.168.1.1/24 start 192.168.1.2 end 192.168.1.254
            exit
            radius-authentication
                password letmein
                radius-server-policy "rad-serv-pol-RSP"
            exit
            radius-proxy-server router 3 name "rad-prox-RPROX"
        exit
    exit
exit
```

RADIUS Proxy Configuration

VPRN-3 is defined for supporting explicit BRG authentication via a RADIUS proxy. The *int-BRG* interface is on SAP 1/1/1:1, and provides connectivity to the management interface of the physical BRGs. The RADIUS proxy listens on the *int-LB-PROXY* interface, and directs the incoming RADIUS messages to the server, as defined by the default authentication-server policy.

```
configure
    service
        vprn 3
            route-distinguisher 64496:3
            interface "int-LB-PROXY" create
                address 10.33.33.1/32
                loopback
            exit
            interface "int-BRG" create
                address 10.1.0.254/24
                sap 1/1/1:1 create
                exit
            exit
            radius-proxy
                server "rad-prox-RPROX" purpose authentication create
                    default-authentication-server-policy "rad-serv-pol-RSP"
                    interface "int-LB-PROXY"
                    secret vsecret1
                    no shutdown
                exit
            exit
            no shutdown
        exit
    exit
exit
```

RADIUS Policies

The RADIUS authentication and accounting policies are defined as follows, so authentication and accounting happens via the base router instance.

```
configure
    router
        radius-server
            server "radius-172.16.1.2" address 172.16.1.2 secret vsecret1 create
                accept-coa
            exit
        exit
    exit
exit

configure
    aaa
```

```
        radius-server-policy "rad-serv-pol-RSP" create
            servers
                router "Base"
                source-address 192.0.2.1
                server 1 name "radius-172.16.1.2"
            exit
        exit
    exit
exit

configure
    subscriber-mgmt
        authentication-policy "radius-AUTH" create
            description "RADIUS authentication policy"
            password letmein
            radius-server-policy "rad-serv-pol-RSP"
        exit
        radius-accounting-policy "radius-ACCT" create
            update-interval 5
            include-radius-attribute
                mac-address
                nat-port-range
                subscriber-id
            exit
            radius-accounting-server
                source-address 192.0.2.1
                router "Base"
                server 1 address 172.16.1.2 secret vsecret1
            exit
        exit
    exit
exit
```

RADIUS User Configuration

Although a PCMP will be used in conjunction with an external controller, for demonstration purposes, this chapter relies on a RADIUS server only.

A sample RADIUS user configuration follows. MAC addresses are used for authentication. MAC address 00:0c:29:00:00:10 identifies the BRG. The addresses ranging from 00:0c:29:00:00:11 to 00:0c:29:00:00:1f identify the home devices connected to that BRG so they all return the same Alc-BRG-Id.

```
00:0c:29:00:00:10       Cleartext-Password := "letmein"
                        Alc-BRG-Id = "00:0c:29:00:00:10",
                        Framed-IPv6-Prefix = 2001:db8:101:1010::/64,
                        Alc-DMZ-address = 192.168.1.254,
                        Alc-Home-Aware-Pool =
                                  "192.168.1.1/24 192.168.1.100-192.168.1.254",
                        Alc-Reserved-Addresses =
                                  "sticky 00:0c:29:00:00:11 192.168.1.110",
                        Alc-Reserved-Addresses +=
                                  "static 00:0c:29:00:00:1f 192.168.1.254",
```

```
                        Alc-Reserved-Addresses +=
                                  "static 00:0c:29:00:00:1e 198.51.100.110",
                        Alc-Portal-Url = "http://11.11.11.11",
                        Alc-Primary-Dns = 1.1.1.1,
                        Alc-Secondary-Dns = 1.1.2.2,
                        Alc-Primary-Nbns = 2.2.1.1,
                        Alc-Secondary-Nbns = 2.2.2.2,
                        Alc-Ipv6-Primary-DNS = 2001:db8:dddd:1::1,
                        Alc-Ipv6-Secondary-DNS = 2001:db8:dddd:2::1,

00:0c:29:00:00:11       Cleartext-Password := "letmein"
                        Alc-BRG-Id = "00:0c:29:00:00:10",

00:0c:29:00:00:12       Cleartext-Password := "letmein"
                        Alc-BRG-Id = "00:0c:29:00:00:10",

00:0c:29:00:00:13       Cleartext-Password := "letmein"
                        Alc-BRG-Id = "00:0c:29:00:00:10",

00:0c:29:00:00:14       Cleartext-Password := "letmein"
                        Alc-BRG-Id = "00:0c:29:00:00:10",
                        Alc-Primary-Dns = 1.1.3.3,
```

Debug Configuration

The following debug configuration can be used for troubleshooting purposes.

```
debug
    router "Base"
        radius
            packet-type authentication accounting coa
            detail-level high
        exit
    exit
    router "1"
        ip
            dhcp
                detail-level medium
                mode egr-ingr-and-dropped
            exit
            icmp6
        exit
    exit
    router "3"
        radius-proxy
            server "rad-prox-RPROX"
                detail-level high
                direction both
                packet-type access-request access-accept access-reject
                            access-challenge accounting-request
                            accounting-response other
            exit
        exit
    exit
exit
```

Explicit Authentication

In the explicit authentication scenario, the BRG is authenticated before any home device attempts to connect. The following trace shows that the RADIUS proxy server in router 3 receives an Access-Request from BRG with BRG ID 00:0c:29:00:00:10, and passes this to the AAA/RADIUS server, which returns the BRG specific data. Two static addresses and one sticky address are associated with this BRG.

```
1 2016/09/20 16:27:22.11 CEST MINOR: DEBUG #2001 vprn3 RADIUS
"RADIUS: Receive
  Proxy-server rad-prox-RPROX"

2 2016/09/20 16:27:22.11 CEST MINOR: DEBUG #2001 vprn3 RADIUS
"RADIUS: Receive
  Access-Request(1) id 18 len 81 from 10.1.0.10:49169 vrid 3 pol rad-serv-pol-RSP
    USER NAME [1] 17 00:0c:29:00:00:10
    PASSWORD [2] 16 bD5mfBsZr5M/aFgOa7iAtE
    NAS IP ADDRESS [4] 4 10.1.0.10
    MESSAGE AUTHENTICATOR [80] 16 0x6c9a951328e303920ba50fa8f7eea0c3
"

3 2016/09/20 16:27:22.11 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.16.1.2:1812 id 1 len 81 vrid 1 pol rad-serv-pol-RSP
    USER NAME [1] 17 00:0c:29:00:00:10
    PASSWORD [2] 16 m9Lo8f7y.V35pw3KzGBs.U
    NAS IP ADDRESS [4] 4 10.1.0.10
    MESSAGE AUTHENTICATOR [80] 16 0xbe54c468e3e6952f0cedfadd3477683b

  Hex Packet Dump:
  --- snipped ---
"

4 2016/09/20 16:27:22.11 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 1 len 388 from 172.16.1.2:1812 vrid 1 pol rad-serv-pol-RSP
    VSA [26] 19 Alcatel(6527)
      BRG ID [225] 17 00:0c:29:00:00:10
    FRAMED IPV6 PREFIX [97] 18 2001:db8:101:1010::/64
    VSA [26] 6 Alcatel(6527)
      BRG DMZ ADDRESS [221] 4 192.168.1.254
    VSA [26] 44 Alcatel(6527)
      BRG HOME AWARE POOL [220] 42 192.168.1.1/24 192.168.1.100-192.168.1.254
    VSA [26] 40 Alcatel(6527)
      BRG RESERVED ADDRESS [223] 38 sticky 00:0c:29:00:00:11 192.168.1.110
    VSA [26] 40 Alcatel(6527)
      BRG RESERVED ADDRESS [223] 38 static 00:0c:29:00:00:1f 192.168.1.254
    VSA [26] 41 Alcatel(6527)
      BRG RESERVED ADDRESS [223] 39 static 00:0c:29:00:00:1e 198.51.100.110
    VSA [26] 20 Alcatel(6527)
      PORTAL URL [177] 18 http://11.11.11.11
    VSA [26] 6 Alcatel(6527)
      PRIMARY DNS [9] 4 1.1.1.1
    VSA [26] 6 Alcatel(6527)
      SECONDARY DNS [10] 4 1.1.2.2
    VSA [26] 6 Alcatel(6527)
```

```
      PRIMARY NBNS [29] 4 2.2.1.1
    VSA [26] 6 Alcatel(6527)
      SECONDARY NBNS [30] 4 2.2.2.2
    VSA [26] 18 Alcatel(6527)
      IPV6 PRIMARY DNS [105] 16 2001:db8:dddd:1::1
    VSA [26] 18 Alcatel(6527)
      IPV6 SECONDARY DNS [106] 16 2001:db8:dddd:2::1

  Hex Packet Dump:
  --- snipped ---
"

5 2016/09/20 16:27:22.12 CEST MINOR: DEBUG #2001 vprn3 RADIUS
"RADIUS: Transmit
  Proxy-server rad-prox-RPROX"

6 2016/09/20 16:27:22.12 CEST MINOR: DEBUG #2001 vprn3 RADIUS
"RADIUS: Transmit
  Access-Accept(2) 10.1.0.10:49169 id 18 len 388 vrid 3
    VSA [26] 19 Alcatel(6527)
      BRG ID [225] 17 00:0c:29:00:00:10
    FRAMED IPV6 PREFIX [97] 18 2001:db8:101:1010::/64
    VSA [26] 6 Alcatel(6527)
      BRG DMZ ADDRESS [221] 4 192.168.1.254
    VSA [26] 44 Alcatel(6527)
      BRG HOME AWARE POOL [220] 42 192.168.1.1/24 192.168.1.100-192.168.1.254
    VSA [26] 40 Alcatel(6527)
      BRG RESERVED ADDRESS [223] 38 sticky 00:0c:29:00:00:11 192.168.1.110
    VSA [26] 40 Alcatel(6527)
      BRG RESERVED ADDRESS [223] 38 static 00:0c:29:00:00:1f 192.168.1.254
    VSA [26] 41 Alcatel(6527)
      BRG RESERVED ADDRESS [223] 39 static 00:0c:29:00:00:1e 198.51.100.110
    VSA [26] 20 Alcatel(6527)
      PORTAL URL [177] 18 http://11.11.11.11
    VSA [26] 6 Alcatel(6527)
      PRIMARY DNS [9] 4 1.1.1.1
    VSA [26] 6 Alcatel(6527)
      SECONDARY DNS [10] 4 1.1.2.2
    VSA [26] 6 Alcatel(6527)
      PRIMARY NBNS [29] 4 2.2.1.1
    VSA [26] 6 Alcatel(6527)
      SECONDARY NBNS [30] 4 2.2.2.2
    VSA [26] 18 Alcatel(6527)
      IPV6 PRIMARY DNS [105] 16 2001:db8:dddd:1::1
    VSA [26] 18 Alcatel(6527)
      IPV6 SECONDARY DNS [106] 16 2001:db8:dddd:2::1
"
```

As a result, the vRGW creates and stores context for this BRG, which can be displayed using the following command. The Proxy authenticated flag is set to “yes”.

```
*A:BNG# show subscriber-mgmt brg gateways

===============================================================================
Bridged Residential Gateways
===============================================================================
Identifier                  : 00:0c:29:00:00:10
SLAAC prefix                : 2001:db8:101:1010::/64
```

```
Subnet                       : 192.168.1.1/24
Subnet start address         : 192.168.1.100
Subnet end address           : 192.168.1.254
DMZ address                  : 192.168.1.254
DNS 1 v4                     : 1.1.1.1
DNS 1 v6                     : 2001:db8:dddd:1::1
DNS 2 v4                     : 1.1.2.2
DNS 2 v6                     : 2001:db8:dddd:2::1
NBNS 1                       : 2.2.1.1
NBNS 2                       : 2.2.2.2
DHCP lease time              : 21600
DHCP stream destination      : (Not Specified)
IPv4 portal URL              : http://11.11.11.11
IPv6 portal URL              : (Not Specified)
BRG profile                  : brg-prof-1
Subscriber profile           : sub-prof-1
SLA profile                  : sla-prof-1
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
===============================================================================
*A:BNG#
```

Initially, no hosts are created and associated with this BRG, as the following command shows.

```
*A:BNG# show subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10" hosts
No entries found.
*A:BNG#
```

To show the static and sticky addresses associated with the BRG, use the following command. Even without any device connected to this BRG, some bindings are created.

```
*A:BNG# show subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10" bindings

===============================================================================
Bridged Residential Gateway home-aware pool address bindings
===============================================================================
Home-aware pool              : 00:0c:29:00:00:10
-------------------------------------------------------------------------------
MAC address                  : 00:0c:29:00:00:11
IP address                   : 192.168.1.110
Allocation type              : sticky-ip-address
DHCP lease                   : false
Remaining lease time         : (Unknown)
Lease start time             : N/A
```

```
MAC address                  : 00:0c:29:00:00:1e
IP address                   : 198.51.100.110
Allocation type              : static
DHCP lease                   : (Unknown)
Remaining lease time         : (Unknown)
Lease start time             : N/A

MAC address                  : 00:0c:29:00:00:1f
IP address                   : 192.168.1.254
Allocation type              : static
DHCP lease                   : (Unknown)
Remaining lease time         : (Unknown)
Lease start time             : N/A

-------------------------------------------------------------------------------
No. of bindings: 3
===============================================================================
*A:BNG#
```

When the first device connects, in this example using DHCPv4 (DORA), this device is authenticated using RADIUS, and the controller returns the corresponding BRG ID and device-specific primary DNS server (messages 7 and 8). Because the BRG has two static addresses associated with it, at the same time these are also authenticated (messages 12 and 15 for the first static host, 13 and 16 for the second static host). The device with MAC address 00:0c:29:00:00:14 has a dedicated primary DNS server, which overrules the primary DNS server defined at BRG level (messages 8, 10, and 14).

```
7 2016/09/20 16:29:08.53 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.16.1.2:1812 id 2 len 79 vrid 1 pol rad-serv-pol-RSP
    USER NAME [1] 17 00:0c:29:00:00:14
    PASSWORD [2] 16 Lh/0pV5SVw5Cp7gBJe75s.
    NAS IP ADDRESS [4] 4 192.0.2.1
    NAS PORT TYPE [61] 4 Ethernet(15)
    NAS PORT ID [87] 8 1/1/1:81

  Hex Packet Dump:
  --- snipped ---
"

8 2016/09/20 16:29:08.53 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 2 len 57 from 172.16.1.2:1812 vrid 1 pol rad-serv-pol-RSP
    VSA [26] 19 Alcatel(6527)
      BRG ID [225] 17 00:0c:29:00:00:10
    VSA [26] 6 Alcatel(6527)
      PRIMARY DNS [9] 4 1.1.3.3

  Hex Packet Dump:
  --- snipped ---
"

9 2016/09/20 16:29:08.53 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
```

```
instance 2 (1), interface index 6 (grp-int-1-1),
   received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:81) Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:0c:29:00:00:14    xid: 0x2

   DHCP options:
   [53] Message type: Discover
   [255] End
"

10 2016/09/20 16:29:08.53 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
   transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:81) Port 68

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 192.168.1.100
   siaddr: 192.168.1.1       giaddr: 192.168.1.1
   chaddr: 00:0c:29:00:00:14    xid: 0x2

   DHCP options:
   [53] Message type: Offer
   [54] DHCP server addr: 192.168.1.1
   [51] Lease time: 21600
   [1] Subnet mask: 255.255.255.0
   [3] Router: 192.168.1.1
   [6] Domain name server: length = 8
            1.1.3.3
            1.1.2.2
   [44] NETBIOS name server: length = 8
            2.2.1.1
            2.2.2.2
   [255] End
"

11 2016/09/20 16:29:08.64 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
   received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:81) Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:0c:29:00:00:14    xid: 0x2

   DHCP options:
   [53] Message type: Request
   [50] Requested IP addr: 192.168.1.100
   [54] DHCP server addr: 192.168.1.1
   [255] End
"

12 2016/09/20 16:29:08.64 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.16.1.2:1812 id 3 len 79 vrid 1 pol rad-serv-pol-RSP
    USER NAME [1] 17 00:0c:29:00:00:1e
```

```
    PASSWORD [2] 16 hW.TR6SdCXXMO/3iZ.3WNk
    NAS IP ADDRESS [4] 4 192.0.2.1
    NAS PORT TYPE [61] 4 Ethernet(15)
    NAS PORT ID [87] 8 1/1/1:81

  Hex Packet Dump:
  --- snipped ---
"

13 2016/09/20 16:29:08.64 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.16.1.2:1812 id 4 len 79 vrid 1 pol rad-serv-pol-RSP
    USER NAME [1] 17 00:0c:29:00:00:1f
    PASSWORD [2] 16 vXTCIxeAZzeGRFAQy8eS/k
    NAS IP ADDRESS [4] 4 192.0.2.1
    NAS PORT TYPE [61] 4 Ethernet(15)
    NAS PORT ID [87] 8 1/1/1:81

  Hex Packet Dump:
  --- snipped ---
"

14 2016/09/20 16:29:08.64 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
   transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:81) Port 68

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 192.168.1.100
   siaddr: 192.168.1.1       giaddr: 192.168.1.1
   chaddr: 00:0c:29:00:00:14    xid: 0x2

   DHCP options:
   [53] Message type: Ack
   [54] DHCP server addr: 192.168.1.1
   [51] Lease time: 21600
   [1] Subnet mask: 255.255.255.0
   [3] Router: 192.168.1.1
   [6] Domain name server: length = 8
            1.1.3.3
            1.1.2.2
   [44] NETBIOS name server: length = 8
            2.2.1.1
            2.2.2.2
   [255] End
"

15 2016/09/20 16:29:08.63 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 3 len 45 from 172.16.1.2:1812 vrid 1 pol rad-serv-pol-RSP
    VSA [26] 19 Alcatel(6527)
      BRG ID [225] 17 00:0c:29:00:00:10

  Hex Packet Dump:
  --- snipped ---
"

16 2016/09/20 16:29:08.63 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
```

```
  Access-Accept(2) id 4 len 45 from 172.16.1.2:1812 vrid 1 pol rad-serv-pol-RSP
    VSA [26] 19 Alcatel(6527)
      BRG ID [225] 17 00:0c:29:00:00:10

  Hex Packet Dump:
  --- snipped ---
"
```

Displaying the bindings again shows that there is now an additional dynamic host, for which the allocation type is dynamic.

```
*A:BNG# show subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10" bindings

===============================================================================
Bridged Residential Gateway home-aware pool address bindings
===============================================================================
Home-aware pool              : 00:0c:29:00:00:10
-------------------------------------------------------------------------------
MAC address                  : 00:0c:29:00:00:11
IP address                   : 192.168.1.110
Allocation type              : sticky-ip-address
DHCP lease                   : false
Remaining lease time         : (Unknown)
Lease start time             : N/A

MAC address                  : 00:0c:29:00:00:14
IP address                   : 192.168.1.100
Allocation type              : dynamic
DHCP lease                   : true
Remaining lease time         : 21580
Lease start time             : 2016/09/20 16:29:08

MAC address                  : 00:0c:29:00:00:1e
IP address                   : 198.51.100.110
Allocation type              : static
DHCP lease                   : (Unknown)
Remaining lease time         : (Unknown)
Lease start time             : N/A

MAC address                  : 00:0c:29:00:00:1f
IP address                   : 192.168.1.254
Allocation type              : static
DHCP lease                   : (Unknown)
Remaining lease time         : (Unknown)
Lease start time             : N/A

-------------------------------------------------------------------------------
No. of bindings: 4
===============================================================================
*A:BNG#
```

The following command shows the hosts associated with this BRG. The sticky address is not in the list because the host is not online. For the sticky address to appear, that device must send a DHCPv4 Discover message, initiating its own authentication.

```
*A:BNG# show subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10" hosts

===============================================================================
Bridged Residential Gateway hosts
===============================================================================
Identifier                  : 00:0c:29:00:00:10
MAC address                 : 00:0c:29:00:00:14
IP address                  : 192.168.1.100
Service                     : 1 (VPRN)
Allocation type             : dynamic
Home-aware pool             : 00:0c:29:00:00:10
DHCP lease                  : true
Remaining lease time        : 21567
Lease start time            : 2016/09/20 16:29:08


Identifier                  : 00:0c:29:00:00:10
MAC address                 : 00:0c:29:00:00:1e
IP address                  : 198.51.100.110
Service                     : 1 (VPRN)
Allocation type             : static
Home-aware pool             : 00:0c:29:00:00:10
DHCP lease                  : (Unknown)
Remaining lease time        : (Unknown)
Lease start time            : N/A


Identifier                  : 00:0c:29:00:00:10
MAC address                 : 00:0c:29:00:00:1f
IP address                  : 192.168.1.254
Service                     : 1 (VPRN)
Allocation type             : static
Home-aware pool             : 00:0c:29:00:00:10
DHCP lease                  : (Unknown)
Remaining lease time        : (Unknown)
Lease start time            : N/A


-------------------------------------------------------------------------------
No. of BRG hosts: 3
===============================================================================
*A:BNG#
```

Even when the last dynamic host disconnects, the static hosts remain, as shown by the following command.

```
*A:BNG# show subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10" hosts

===============================================================================
Bridged Residential Gateway hosts
===============================================================================
Identifier                  : 00:0c:29:00:00:10
MAC address                 : 00:0c:29:00:00:1e
IP address                  : 198.51.100.110
Service                     : 1 (VPRN)
Allocation type             : static
Home-aware pool             : 00:0c:29:00:00:10
DHCP lease                  : (Unknown)
```

```
Remaining lease time         : (Unknown)
Lease start time             : N/A


Identifier                   : 00:0c:29:00:00:10
MAC address                  : 00:0c:29:00:00:1f
IP address                   : 192.168.1.254
Service                      : 1 (VPRN)
Allocation type              : static
Home-aware pool              : 00:0c:29:00:00:10
DHCP lease                   : (Unknown)
Remaining lease time         : (Unknown)
Lease start time             : N/A


-------------------------------------------------------------------------------
No. of BRG hosts: 2
===============================================================================
*A:BNG#
```

As long as the BRG is still alive, these hosts will remain, and so will the BRG context. For that reason, the vRGW might start connectivity verification and eventually delete the BRG and its hosts, depending on the configuration.

The BRG context can also be cleared manually using the following command, after which the BRG and the associated hosts are deleted.

```
*A:BNG# clear subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10"

*A:BNG# show subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10" hosts
No entries found.

*A:BNG# show subscriber-mgmt brg gateways
No entries found.
*A:BNG#
```

Implicit Authentication

In the implicit authentication scenario, the BRG is authenticated when the first host is authenticated, which requires two phases. In the first phase, the RADIUS server is accessed for authenticating the device (00:0c:29:00:00:11), returning the device-specific data including the BRG ID. In the second phase, the RADIUS server is accessed for authenticating the BRG (00:0c:29:00:00:10), returning the home-specific data. Because the RADIUS server returns two reserved static addresses, SR OS additionally authenticates the static devices.

```
18 2016/09/20 16:30:59.88 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.16.1.2:1812 id 5 len 79 vrid 1 pol rad-serv-pol-RSP
    USER NAME [1] 17 00:0c:29:00:00:11
    PASSWORD [2] 16 OhUXqlDfL7XPKlx8EmoMak
```

```
    NAS IP ADDRESS [4] 4 192.0.2.1
    NAS PORT TYPE [61] 4 Ethernet(15)
    NAS PORT ID [87] 8 1/1/1:81

  Hex Packet Dump:
  --- snipped ---
"

19 2016/09/20 16:30:59.89 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 5 len 45 from 172.16.1.2:1812 vrid 1 pol rad-serv-pol-RSP
    VSA [26] 19 Alcatel(6527)
      BRG ID [225] 17 00:0c:29:00:00:10

  Hex Packet Dump:
  --- snipped ---
"

20 2016/09/20 16:30:59.89 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.16.1.2:1812 id 6 len 88 vrid 1 pol rad-serv-pol-RSP
    USER NAME [1] 17 00:0c:29:00:00:10
    PASSWORD [2] 16 u87pDAyiQx.TVIGEHFV22U
    NAS IP ADDRESS [4] 4 192.0.2.1
    VSA [26] 19 Alcatel(6527)
      BRG ID [225] 17 00:0c:29:00:00:10

  Hex Packet Dump:
  --- snipped ---
"

21 2016/09/20 16:30:59.89 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 6 len 388 from 172.16.1.2:1812 vrid 1 pol rad-serv-pol-RSP
    VSA [26] 19 Alcatel(6527)
      BRG ID [225] 17 00:0c:29:00:00:10
    FRAMED IPV6 PREFIX [97] 18 2001:db8:101:1010::/64
    VSA [26] 6 Alcatel(6527)
      BRG DMZ ADDRESS [221] 4 192.168.1.254
    VSA [26] 44 Alcatel(6527)
      BRG HOME AWARE POOL [220] 42 192.168.1.1/24 192.168.1.100-192.168.1.254
    VSA [26] 40 Alcatel(6527)
      BRG RESERVED ADDRESS [223] 38 sticky 00:0c:29:00:00:11 192.168.1.110
    VSA [26] 40 Alcatel(6527)
      BRG RESERVED ADDRESS [223] 38 static 00:0c:29:00:00:1f 192.168.1.254
    VSA [26] 41 Alcatel(6527)
      BRG RESERVED ADDRESS [223] 39 static 00:0c:29:00:00:1e 198.51.100.110
    VSA [26] 20 Alcatel(6527)
      PORTAL URL [177] 18 http://11.11.11.11
    VSA [26] 6 Alcatel(6527)
      PRIMARY DNS [9] 4 1.1.1.1
    VSA [26] 6 Alcatel(6527)
      SECONDARY DNS [10] 4 1.1.2.2
    VSA [26] 6 Alcatel(6527)
      PRIMARY NBNS [29] 4 2.2.1.1
    VSA [26] 6 Alcatel(6527)
      SECONDARY NBNS [30] 4 2.2.2.2
    VSA [26] 18 Alcatel(6527)
      IPV6 PRIMARY DNS [105] 16 2001:db8:dddd:1::1
```

```
    VSA [26] 18 Alcatel(6527)
      IPV6 SECONDARY DNS [106] 16 2001:db8:dddd:2::1

  Hex Packet Dump:
  --- snipped ---
"

22 2016/09/20 16:30:59.89 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
   received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:81) Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:0c:29:00:00:11    xid: 0x1

   DHCP options:
   [53] Message type: Discover
   [255] End
"

23 2016/09/20 16:30:59.89 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
   transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:81) Port 68

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 192.168.1.110
   siaddr: 192.168.1.1       giaddr: 192.168.1.1
   chaddr: 00:0c:29:00:00:11    xid: 0x1

   DHCP options:
   [53] Message type: Offer
   [54] DHCP server addr: 192.168.1.1
   [51] Lease time: 21600
   [1] Subnet mask: 255.255.255.0
   [3] Router: 192.168.1.1
   [6] Domain name server: length = 8
            1.1.1.1
            1.1.2.2
   [44] NETBIOS name server: length = 8
            2.2.1.1
            2.2.2.2
   [255] End
"

24 2016/09/20 16:31:00.00 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
   received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:81) Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 00:0c:29:00:00:11    xid: 0x1

   DHCP options:
   [53] Message type: Request
```

```
   [50] Requested IP addr: 192.168.1.110
   [54] DHCP server addr: 192.168.1.1
   [255] End
"

25 2016/09/20 16:31:00.00 CEST MINOR: DEBUG #2001 vprn1 PIP
"PIP: DHCP
instance 2 (1), interface index 6 (grp-int-1-1),
   transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:81) Port 68

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 192.168.1.110
   siaddr: 192.168.1.1       giaddr: 192.168.1.1
   chaddr: 00:0c:29:00:00:11    xid: 0x1

   DHCP options:
   [53] Message type: Ack
   [54] DHCP server addr: 192.168.1.1
   [51] Lease time: 21600
   [1] Subnet mask: 255.255.255.0
   [3] Router: 192.168.1.1
   [6] Domain name server: length = 8
            1.1.1.1
            1.1.2.2
   [44] NETBIOS name server: length = 8
            2.2.1.1
            2.2.2.2
   [255] End
"

26 2016/09/20 16:31:00.00 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.16.1.2:1812 id 7 len 79 vrid 1 pol rad-serv-pol-RSP
    USER NAME [1] 17 00:0c:29:00:00:1e
    PASSWORD [2] 16 hz9D.7GnC2H.LSaqMZ4TiE
    NAS IP ADDRESS [4] 4 192.0.2.1
    NAS PORT TYPE [61] 4 Ethernet(15)
    NAS PORT ID [87] 8 1/1/1:81

  Hex Packet Dump:
  --- snipped ---
"

27 2016/09/20 16:31:00.00 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 172.16.1.2:1812 id 8 len 79 vrid 1 pol rad-serv-pol-RSP
    USER NAME [1] 17 00:0c:29:00:00:1f
    PASSWORD [2] 16 zIWsyfz7KatfZ42IB4/uY.
    NAS IP ADDRESS [4] 4 192.0.2.1
    NAS PORT TYPE [61] 4 Ethernet(15)
    NAS PORT ID [87] 8 1/1/1:81

  Hex Packet Dump:
  --- snipped ---
"

28 2016/09/20 16:30:59.99 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 7 len 45 from 172.16.1.2:1812 vrid 1 pol rad-serv-pol-RSP
```

```
    VSA [26] 19 Alcatel(6527)
      BRG ID [225] 17 00:0c:29:00:00:10

  Hex Packet Dump:
  --- snipped ---
"

29 2016/09/20 16:30:59.99 CEST MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 8 len 45 from 172.16.1.2:1812 vrid 1 pol rad-serv-pol-RSP
    VSA [26] 19 Alcatel(6527)
      BRG ID [225] 17 00:0c:29:00:00:10

  Hex Packet Dump:
  --- snipped ---
"
```

Subsequent connections of additional devices connected to the BRG result in one single RADIUS access per device. There is no need to reauthenticate the BRG.no sched

```
30 2016/09/20 16:32:10.74 CEST MINOR: DEBUG #2001 vprn1 PIPno sched
"PIP: DHCPno sched
instance 2 (1), interface index 6 (grp-int-1-1), no sched
   received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:81) Port 67no sched
no sched
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6no sched
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0no sched
   siaddr: 0.0.0.0           giaddr: 0.0.0.0no sched
   chaddr: 00:0c:29:00:00:14    xid: 0x2no sched
no sched
   DHCP options:no sched
   [53] Message type: Discoverno sched
   [255] Endno sched
"no sched
no sched
31 2016/09/20 16:32:10.74 CEST MINOR: DEBUG #2001 Base RADIUSno sched
"RADIUS: Transmitno sched
  Access-Request(1) 172.16.1.2:1812 id 9 len 79 vrid 1 pol rad-serv-pol-RSPno sched
    USER NAME [1] 17 00:0c:29:00:00:14no sched
    PASSWORD [2] 16 TcHWkAooYb5Tcpj9KPR3M.no sched
    NAS IP ADDRESS [4] 4 192.0.2.1no sched
    NAS PORT TYPE [61] 4 Ethernet(15)no sched
    NAS PORT ID [87] 8 1/1/1:81no sched
no sched
  Hex Packet Dump:no sched
  --- snipped ---no sched
"no sched
no sched
32 2016/09/20 16:32:10.74 CEST MINOR: DEBUG #2001 Base RADIUSno sched
"RADIUS: Receiveno sched
  Access-Accept(2) id 9 len 57 from 172.16.1.2:1812 vrid 1 pol rad-serv-pol-RSPno
sched
    VSA [26] 19 Alcatel(6527)no sched
      BRG ID [225] 17 00:0c:29:00:00:10no sched
    VSA [26] 6 Alcatel(6527)no sched
      PRIMARY DNS [9] 4 1.1.3.3no sched
no sched
```

```
  Hex Packet Dump:no sched
  --- snipped ---no sched
"no sched
no sched
33 2016/09/20 16:32:10.74 CEST MINOR: DEBUG #2001 vprn1 PIPno sched
"PIP: DHCPno sched
instance 2 (1), interface index 6 (grp-int-1-1), no sched
   transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:81) Port 68no sched
no sched
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6no sched
   ciaddr: 0.0.0.0           yiaddr: 192.168.1.100no sched
   siaddr: 192.168.1.1       giaddr: 192.168.1.1no sched
   chaddr: 00:0c:29:00:00:14    xid: 0x2no sched
no sched
   DHCP options:no sched
   [53] Message type: Offerno sched
   [54] DHCP server addr: 192.168.1.1no sched
   [51] Lease time: 21600no sched
   [1] Subnet mask: 255.255.255.0no sched
   [3] Router: 192.168.1.1no sched
   [6] Domain name server: length = 8no sched
            1.1.3.3no sched
            1.1.2.2no sched
   [44] NETBIOS name server: length = 8no sched
            2.2.1.1no sched
            2.2.2.2no sched
   [255] Endno sched
"no sched
no sched
34 2016/09/20 16:32:10.84 CEST MINOR: DEBUG #2001 vprn1 PIPno sched
"PIP: DHCPno sched
instance 2 (1), interface index 6 (grp-int-1-1), no sched
   received DHCP Boot Request on Interface grp-int-1-1 (1/1/1:81) Port 67no sched
no sched
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6no sched
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0no sched
   siaddr: 0.0.0.0           giaddr: 0.0.0.0no sched
   chaddr: 00:0c:29:00:00:14    xid: 0x2no sched
no sched
   DHCP options:no sched
   [53] Message type: Requestno sched
   [50] Requested IP addr: 192.168.1.100no sched
   [54] DHCP server addr: 192.168.1.1no sched
   [255] Endno sched
"no sched
no sched
35 2016/09/20 16:32:10.84 CEST MINOR: DEBUG #2001 vprn1 PIPno sched
"PIP: DHCPno sched
instance 2 (1), interface index 6 (grp-int-1-1), no sched
   transmitted DHCP Boot Reply to Interface grp-int-1-1 (1/1/1:81) Port 68no sched
no sched
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6no sched
   ciaddr: 0.0.0.0           yiaddr: 192.168.1.100no sched
   siaddr: 192.168.1.1       giaddr: 192.168.1.1no sched
   chaddr: 00:0c:29:00:00:14    xid: 0x2no sched
no sched
   DHCP options:no sched
   [53] Message type: Ackno sched
   [54] DHCP server addr: 192.168.1.1no sched
```

```
   [51] Lease time: 21600no sched
   [1] Subnet mask: 255.255.255.0no sched
   [3] Router: 192.168.1.1no sched
   [6] Domain name server: length = 8no sched
            1.1.3.3no sched
            1.1.2.2no sched
   [44] NETBIOS name server: length = 8no sched
            2.2.1.1no sched
            2.2.2.2no sched
   [255] Endno sched
"no sched
```

The following command shows the corresponding BRG hosts.

```
*A:BNG# show subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10" hosts

===============================================================================
Bridged Residential Gateway hosts
===============================================================================
Identifier                  : 00:0c:29:00:00:10
MAC address                 : 00:0c:29:00:00:11
IP address                  : 192.168.1.110
Service                     : 1 (VPRN)
Allocation type             : sticky-ip-address
Home-aware pool             : 00:0c:29:00:00:10
DHCP lease                  : true
Remaining lease time        : 21514
Lease start time            : 2016/09/20 16:31:00


Identifier                  : 00:0c:29:00:00:10
MAC address                 : 00:0c:29:00:00:14
IP address                  : 192.168.1.100
Service                     : 1 (VPRN)
Allocation type             : dynamic
Home-aware pool             : 00:0c:29:00:00:10
DHCP lease                  : true
Remaining lease time        : 21584
Lease start time            : 2016/09/20 16:32:10


Identifier                  : 00:0c:29:00:00:10
MAC address                 : 00:0c:29:00:00:1e
IP address                  : 198.51.100.110
Service                     : 1 (VPRN)
Allocation type             : static
Home-aware pool             : 00:0c:29:00:00:10
DHCP lease                  : (Unknown)
Remaining lease time        : (Unknown)
Lease start time            : N/A


Identifier                  : 00:0c:29:00:00:10
MAC address                 : 00:0c:29:00:00:1f
IP address                  : 192.168.1.254
Service                     : 1 (VPRN)
Allocation type             : static
Home-aware pool             : 00:0c:29:00:00:10
```

```
DHCP lease                   : (Unknown)
Remaining lease time         : (Unknown)
Lease start time             : N/A


-------------------------------------------------------------------------------
No. of BRG hosts: 4
===============================================================================
*A:BNG#
```

Releasing both dynamic devices results in all BRG hosts being deleted, including the static hosts.

```
*A:BNG# show subscriber-mgmt brg gateway brg-id "00:0c:29:00:00:10" hosts
No entries found.
*A:BNG#
```

Also, the BRG is deleted automatically.

```
*A:BNG# show subscriber-mgmt brg gateways
No entries found.
*A:BNG#
```

Conclusion

This chapter describes the explicit and the implicit authentication models for the BRG. The explicit authentication model requires the BRG to contain an embedded RADIUS client, and offers better security in comparison with the implicit model, which does not require an embedded RADIUS client.

Virtual Residential Gateway Home Pool Management

This chapter describes virtual residential gateway home pool management.

Topics in this chapter include:

Applicability

This chapter is applicable to SR OS routers and is based on SR OS 14.0.R3.

Overview

In the virtual residential gateway model, the Layer 3 (L3) functions are moved out of the traditional residential gateway (RGW) and into the network; see Figure 224. Examples of L3 functions moved to the network are:

- DHCPv4
- Network Address Translation (NAT)
- Firewalling
- Universal Plug and Play

The in-home equipment interconnecting all devices in the home is referred to as the bridged residential gateway (BRG). The BRG only handles Layer 2 (L2) connectivity (for example, Ethernet and WiFi) and is always operating in bridged mode. The BRG has an L2 uplink connecting it to the virtual residential gateway (vRGW), either through a direct link or through tunneling technology. The vRGW handles all L3 connectivity.

For the vRGW to offer IPv4 connectivity to the devices in the home network, the vRGW provides the following features:

- Private addresses from a single home address pool are offered. The address pools can overlap between homes.
- Sticky or static addresses provide a fixed device to IP mapping.
- Public addresses can be assigned, to enable home servers to be publicly accessible.
- A demilitarized (DMZ) host can be defined.

Because multiple homes are allowed to use the same private subnet, the vRGW requires L2-aware NAT. L2-aware NAT is handled in the NAT in Combination with ESMchapter.

For the vRGW to offer IPv6 connectivity to the devices in the home network, the vRGW provides the following features:

- A /64 SLAAC prefix is assigned per home.
- IA_NA address allocation using DHCPv6 relay or proxy is supported, following standard ESM rules.
- Prefix Delegation (PD) is not supported.

Using the vRGW model, the ISP now has visibility on the MAC and IP addresses used by the in-home devices.

Figure 224 **Virtual Residential Gateway in the Network with Bridged Residential Gateway at Home**

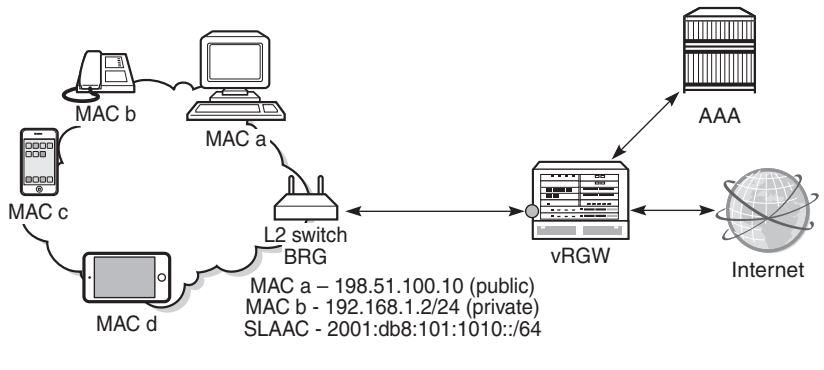

DHCP and IP Address Management

The vRGW has the following characteristics for DHCPv4:

- One pool per home
- IP overlap between homes
- Sticky IP addresses

Sticky IP addresses are DHCPv4 addresses assigned to devices that need to have the same address all the time and are provided through the DHCP protocol. Home servers, network-attached storage (NAS), and network printers are examples of devices that typically are configured with sticky IP addresses. The vRGW sets a flag indicating that the IP address is reserved, to avoid assigning the sticky address to devices that do not have this requirement.

Static IP addresses can be used and configured for devices that do not use or support the DHCP protocol, and are configured manually on the device. The static IP address used can be a public or a private address. Traffic to and from private addresses undergoes NAT, whereas traffic to and from public addresses does not undergo NAT. The vRGW drops DHCP messages originating from devices that are considered to be static.

The vRGW requires the use of IPoE sessions for supporting BRGs, and creates an IPoE session per device. Static devices require at least one dynamic host to be created first so that the associated SAP or tunnel can be defined to send traffic over. Each static device is authenticated individually using RADIUS to retrieve the per-device parameters.

The vRGW has the following characteristics for IPv6:

- Both SLAAC and IA_NA allocation are supported.
- Static IPv6 addresses are not supported. However, a static IPv4 device can get an IPv6 SLAAC prefix if IPoE-linking is enabled.
- No per-home pool is supported for IA_NA, but DHCPv6 can be relayed to an external or local DHCPv6 server.

The vRGW only operates in IPoE bridged mode, so that multiple hosts on the same BRG share the same SLAAC/64 prefix. Only one SLAAC prefix per BRG is supported.

Prefix delegation (PD) is not supported in the vRGW.

Demilitarized Zone (DMZ)

When a DMZ host is specified within a home, all downstream traffic not matching an existing flow or port forwarding rule is sent to that host. Without a DMZ host, this traffic is dropped. DMZ requires a single port-reservation block to be defined on the outside L2-aware NAT range.

Configuration

The services configuration overview shown in Figure 225 applies to the examples used throughout this chapter. For L2-aware NAT, two routed services are configured. In this example, VPRN-1 provides the connection toward the customer premises, and hosts the NAT inside addresses. VPRN-2 provides the connection toward the Internet, and hosts the NAT outside addresses. Addresses from the documentation range as defined in RFC 5737 are used for the public addresses.

The examples in this chapter use capture and managed SAPs. The capture SAP is created manually on VPLS-10, and the managed SAPs are created dynamically on VPRN-1. Also, static SAPs can be used to support BRGs, but in that case the configuration file can grow rapidly when many SAPs are required.

Figure 225 **Services Configuration Overview**

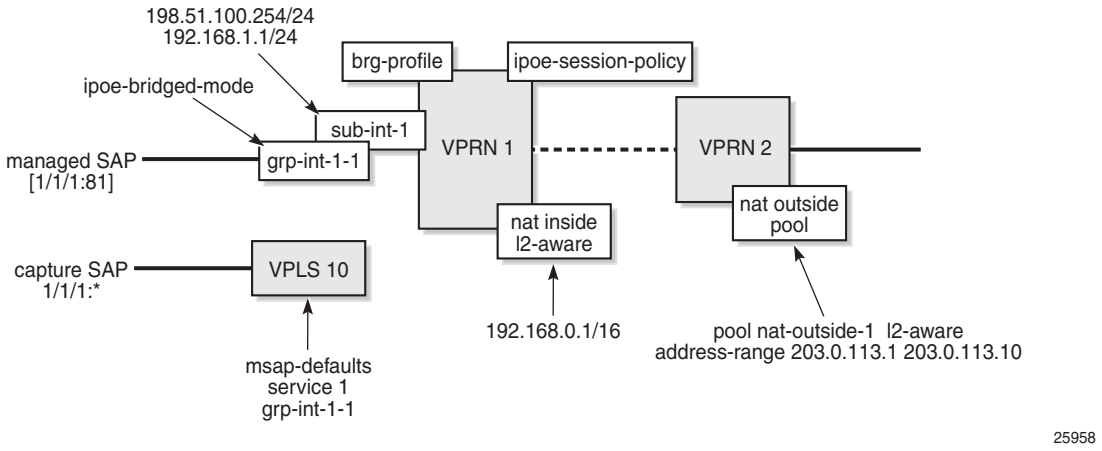

Services Configuration

An excerpt of the configuration of VPRN-1 follows. This VPRN contains subscriber interface *sub-int-1* with group interface *grp-int-1-1*, thereby using the routed central office (CO) model. No static SAPs are configured; managed SAPs will be created when triggers are received on VPLS-10. The IPoE session policy required for the vRGW is *sess-pol-SAP-MAC*, and the default BRG profile used is *brg-prof-1*.

Relay and proxy scenarios are configured for DHCPv4 and DHCPv6. SLAAC prefixes are taken from the wan-host prefix range, and must be advertised to devices using router-advertisement (RA) messages in response to router-solicit (RS) messages. For that purpose, group interface *grp-int-1-1* is configured in its IPv6 context to support router advertisement and router solicitation. Also IPoE bridged mode is enabled in this context, so that the same SLAAC prefix can be allocated multiple times to the same SAP.

RADIUS authentication is enabled through authentication policy *radius-AUTH*. The inside L2-aware NAT range is 192.168.0.1/16, and the DHCP pool subnets defined in the BRG profile must belong to that range; see the ESM Configuration part in this chapter.

```
configure
    service
        vprn 1 customer 1 create
            --- snipped ---
            route-distinguisher 64496:1
            interface "int-DHCP" create
                address 10.11.11.1/32
                ipv6
                    address 2001:db8::11/128
                    local-dhcp-server "dhcp6-SRVC1"
                exit
                loopback
            exit
            subscriber-interface "sub-int-1" create
                address 198.51.100.254/24
                address 192.168.1.1/24
                ipv6
                    link-local-address fe80::1
                    subscriber-prefixes
                        prefix 2001:db8:101::/48 wan-host
                    exit
                exit
                group-interface "grp-int-1-1" create
                    ipv6
                        router-advertisements
                            prefix-options
                                autonomous
                            exit
                            no shutdown
                        exit
                        dhcp6
                            proxy-server
```

```
                                    no shutdown
                                exit
                                relay
                                    server 2001:db8::11
                                    no shutdown
                                exit
                            exit
                            router-solicit
                                inactivity-timer hrs 2
                                no shutdown
                            exit
                            ipoe-bridged-mode
                        exit
                        --- snipped ---
                        arp-populate
                        dhcp
                            proxy-server
                                emulated-server 198.51.100.254
                                no shutdown
                            exit
                            trusted
                            lease-populate 128
                            gi-address 198.51.100.254
                            no shutdown
                        exit
                        authentication-policy "radius-AUTH"
                        ipoe-session
                            ipoe-session-policy "sess-pol-SAP-MAC"
                            sap-session-limit 128
                            no shutdown
                        exit
                        brg
                            default-brg-profile "brg-prof-1"
                            no shutdown
                        exit
                        oper-up-while-empty
                    exit
                exit
                nat
                    inside
                        l2-aware
                            address 192.168.0.1/16
                        exit
                    exit
                exit
                no shutdown
            exit
        exit
exit
```

VPRN-2 defines an interface to the outside world, *int-VPRN2-INTERNET*, as well as the outside L2-aware NAT range, which is 203.0.113.1 up to 203.0.113.10. Port-reservation blocks is set to 1, to ensure a unique outside IP per subscriber (home) and correct operation of the DMZ feature.

```
configure
    service
        vprn 2 customer 1 create
            route-distinguisher 64496:2
            interface "int-VPRN2-INTERNET" create
                --- snipped ---
            exit
            nat
                outside
                    pool "nat-outside-1" nat-group 1 type l2-aware create
                        port-reservation blocks 1
                        address-range 203.0.113.1 203.0.113.10 create
                        exit
                        no shutdown
                    exit
                exit
            exit
            no shutdown
        exit
    exit
exit
```

VPLS-10 defines the capture SAP on port 1/1/1. The triggers configured are dhcp, dhcp6, and rtr-solicit. The authentication policy and the IPoE session policy used are the same as the ones used on VPRN-1, and are *radius-AUTH* and *sess-pol-SAP-MAC*, respectively. The MSAP defaults indicate that the managed SAPs must be created on service 1, group interface *grp-int-1-1*, using MSAP policy *msap-pol-MSAP*.

```
configure
    service
        vpls 10 customer 1 create
            description "VPLS for capture SAPs - BRG-demo"
            stp
                shutdown
            exit
            sap 1/1/1:* capture-sap create
                description "capture SAP for MSAP creation on 1/1/1"
                trigger-packet dhcp dhcp6 rtr-solicit
                ipoe-session
                    ipoe-session-policy "sess-pol-SAP-MAC"
                    no shutdown
                exit
                msap-defaults
                    group-interface "grp-int-1-1"
                    policy "msap-pol-MSAP"
                    service 1
                exit
                authentication-policy "radius-AUTH"
                no shutdown
            exit
            no shutdown
        exit
    exit
exit
```

RADIUS User Configuration

Although a protocol conversion and mediation platform (PCMP) will be used in conjunction with an external controller, for demonstration purposes this chapter relies on a RADIUS server only.

A sample RADIUS user configuration follows. The user's MAC address is used for authenticating purposes. MAC address 00:0c:29:00:00:10 identifies the BRG. The addresses ranging from 00:0c:29:00:00:11 to 00:0c:29:00:00:1f identify the home devices connected to the same BRG so they all return the same Alc-BRG-Id. For the BRG, RADIUS provides primary and secondary DNS and NBNS servers, a set of reserved addresses, a DMZ address, a framed IPv6 prefix (used for SLAAC), a home-aware pool, and the BRG profile. The home pool subnet must be a subnet of the L2-aware inside subnet.

```
00:0c:29:00:00:10       Cleartext-Password := "letmein"
                        Alc-BRG-Id = "00:0c:29:00:00:10",
                        Alc-BRG-Profile = "brg-prof-1",
                        Framed-IPv6-Prefix = 2001:db8:101:1010::/64,
                        Alc-DMZ-address = 192.168.1.254,
                        Alc-Home-Aware-Pool =
                                   "192.168.1.1/24 192.168.1.100-192.168.1.254",
                        Alc-Reserved-Addresses =
                                   "sticky 00:0c:29:00:00:11 192.168.1.110",
                        Alc-Reserved-Addresses +=
                                   "static 00:0c:29:00:00:1f 192.168.1.254",
                        Alc-Reserved-Addresses +=
                                   "static 00:0c:29:00:00:1e 198.51.100.110",
                        Alc-Portal-Url = "http://11.11.11.11",
                        Alc-Primary-Dns = 1.1.1.1,
                        Alc-Secondary-Dns = 1.1.2.2,
                        Alc-Primary-Nbns = 2.2.1.1,
                        Alc-Secondary-Nbns = 2.2.2.2,
                        Alc-Ipv6-Primary-DNS = 2001:db8:dddd:1::1,
                        Alc-Ipv6-Secondary-DNS = 2001:db8:dddd:2::1,

00:0c:29:00:00:11       Cleartext-Password := "letmein"
                        Alc-BRG-Id = "00:0c:29:00:00:10",

00:0c:29:00:00:12       Cleartext-Password := "letmein"
                        Alc-BRG-Id = "00:0c:29:00:00:10",

00:0c:29:00:00:13       Cleartext-Password := "letmein"
                        Alc-BRG-Id = "00:0c:29:00:00:10",

00:0c:29:00:00:14       Cleartext-Password := "letmein"
                        Alc-BRG-Id = "00:0c:29:00:00:10",
                        Alc-Primary-Dns = 1.1.3.3,
--- snipped ---
```

The RADIUS server is allowed to return a SLAAC pool name using the *Alc-SLAAC-IPv6-Pool* attribute, but in that case local address assignment needs to be configured on the vRGW. See the ESM SLAAC Prefix Assignment via Local Address Server chapter for more information.

ESM Configuration

The subscriber management policies and profiles are defined as follows. The BRG profile *brg-prof-1* provides an SLA profile, a sub-profile, a RADIUS server policy, plus the corresponding password, a RADIUS proxy server, and a DHCP pool. The DHCP pool is within the inside L2-aware NAT range defined for VPRN-1.

```
configure
    subscriber-mgmt
        ipoe-session-policy "sess-pol-SAP-MAC" create
        exit
        sla-profile "sla-prof-1" create
        exit
        sub-profile "sub-prof-1" create
            nat-policy "nat-pol-1"
        exit
        sub-ident-policy "sub-ident-DIRECT" create
            sub-profile-map
                use-direct-map-as-default
            exit
            sla-profile-map
                use-direct-map-as-default
            exit
        exit
        msap-policy "msap-pol-MSAP" create
            sub-sla-mgmt
                def-sub-id use-sap-id
                def-sub-profile "sub-prof-1"
                def-sla-profile "sla-prof-1"
                sub-ident-policy "sub-ident-DIRECT"
            exit
        exit
        brg-profile "brg-prof-1" create
            description "default BRG-profile, demo purposes"
            sla-profile-string "sla-prof-1"
            sub-profile-string "sub-prof-1"
            dhcp-pool
                subnet 192.168.1.1/24 start 192.168.1.2 end 192.168.1.254
            exit
            radius-authentication
                password letmein
                radius-server-policy "rad-serv-pol-RSP"
            exit
        exit
    exit
exit
```

The RADIUS authentication and accounting policies are defined as follows.

```
configure
    router
        radius-server
            server "radius-172.16.1.2" address 172.16.1.2 secret vsecret1 create
                accept-coa
            exit
        exit
    exit
exit

configure
    aaa
        radius-server-policy "rad-serv-pol-RSP" create
            servers
                router "Base"
                source-address 192.0.2.1
                server 1 name "radius-172.16.1.2"
            exit
        exit
    exit
exit

configure
    subscriber-mgmt
        authentication-policy "radius-AUTH" create
            description "Radius authentication policy"
            password letmein
            radius-authentication-server
                source-address 192.0.2.1
            exit
            radius-server-policy "rad-serv-pol-RSP"
        exit
        radius-accounting-policy "radius-ACCT" create
            update-interval 5
            include-radius-attribute
                mac-address
                nat-port-range
                subscriber-id
            exit
            radius-accounting-server
                source-address 192.0.2.1
                router "Base"
                server 1 address 172.16.1.2 secret vsecret1
            exit
        exit
    exit
exit
```

NAT Policy configuration

The NAT policy used in support for the BRGs is *nat-pol-1*, and refers to the outside address pool defined in VPRN-2.

```
configure
    service
        nat
            nat-policy "nat-pol-1" create
                pool "nat-outside-1" router 2
            exit
        exit
    exit
exit
```

Operation and Verification

The following command shows the current BRG hosts. Six hosts are connected. The first host has a sticky address, the third a plain dynamic address, the fifth a public static address, and the last a private static address. The second and the fourth hosts correspond to the SLAAC hosts, and their allocation type is “not-applicable”. For the static hosts, no DHCP lease information is maintained.

```
*A:BNG# show subscriber-mgmt brg brg-hosts

===============================================================================
Bridged Residential Gateway hosts
===============================================================================
Identifier                  : 00:0c:29:00:00:10
MAC address                 : 00:0c:29:00:00:11
IP address                  : 192.168.1.110
Service                     : 1 (VPRN)
Allocation type             : sticky-ip-address
Home-aware pool             : 00:0c:29:00:00:10
DHCP lease                  : true
Remaining lease time        : 11479
Lease start time            : 2016/09/19 20:51:22


Identifier                  : 00:0c:29:00:00:10
MAC address                 : 00:0c:29:00:00:11
IP address                  : 2001:db8:101:1010::
Service                     : 1 (VPRN)
Allocation type             : not-applicable


Identifier                  : 00:0c:29:00:00:10
MAC address                 : 00:0c:29:00:00:14
IP address                  : 192.168.1.100
Service                     : 1 (VPRN)
Allocation type             : dynamic
Home-aware pool             : 00:0c:29:00:00:10
DHCP lease                  : true
Remaining lease time        : 11480
Lease start time            : 2016/09/19 20:51:24


Identifier                  : 00:0c:29:00:00:10
MAC address                 : 00:0c:29:00:00:14
```

```
IP address                  : 2001:db8:101:1010::
Service                     : 1 (VPRN)
Allocation type             : not-applicable


Identifier                  : 00:0c:29:00:00:10
MAC address                 : 00:0c:29:00:00:1e
IP address                  : 198.51.100.110
Service                     : 1 (VPRN)
Allocation type             : static
Home-aware pool             : 00:0c:29:00:00:10
DHCP lease                  : (Unknown)
Remaining lease time        : (Unknown)
Lease start time            : N/A


Identifier                  : 00:0c:29:00:00:10
MAC address                 : 00:0c:29:00:00:1f
IP address                  : 192.168.1.254
Service                     : 1 (VPRN)
Allocation type             : static
Home-aware pool             : 00:0c:29:00:00:10
DHCP lease                  : (Unknown)
Remaining lease time        : (Unknown)
Lease start time            : N/A


-------------------------------------------------------------------------------
No. of BRG hosts: 6
===============================================================================
*A:BNG#
```

The following command shows the active BRGs. Because all the hosts from the previous command belong to the same BRG, only a single BRG gateway exists. The SLAAC prefix, subnet, start address, end address, DMZ address, DNS addresses for IPv4 and IPv6, NBNS-1 and NBNS-2 addresses, and IPv4 and IPv6 portal addresses are obtained from the RADIUS server. “DMZ address in use” is set to “yes” because RADIUS returned the Alc-DMZ-address, and the outside L2 NAT pool has a single port-reservation block configured (ref. VPRN-2).

```
*A:BNG# show subscriber-mgmt brg gateways

===============================================================================
Bridged Residential Gateways
===============================================================================
Identifier                  : 00:0c:29:00:00:10
SLAAC prefix                : 2001:db8:101:1010::/64
Subnet                      : 192.168.1.1/24
Subnet start address        : 192.168.1.100
Subnet end address          : 192.168.1.254
DMZ address                 : 192.168.1.254
DNS 1 v4                    : 1.1.1.1
DNS 1 v6                    : 2001:db8:dddd:1::1
DNS 2 v4                    : 1.1.2.2
DNS 2 v6                    : 2001:db8:dddd:2::1
NBNS 1                      : 2.2.1.1
```

```
NBNS 2                        : 2.2.2.2
DHCP lease time               : 21600
DHCP stream destination       : (Not Specified)
IPv4 portal URL               : http://11.11.11.11
IPv6 portal URL               : (Not Specified)
BRG profile                   : brg-prof-1
Subscriber profile            : sub-prof-1
SLA profile                   : sla-prof-1
UPnP policy override          : (Not Specified)
DMZ address in use            : yes
Proxy authenticated           : no
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
===============================================================================
*A:BNG#
```

The following command shows the active subscribers. Only a single subscriber exists, with two SLAAC hosts (origin is SLAAC), two dynamic hosts (origin is DHCP), and two “static” hosts (origin is AAA). Also, the NAT policy used and the outside IP address and the ports are shown.

```
*A:BNG# show service active-subscribers
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber 00:0c:29:00:00:10 (sub-prof-1)
-------------------------------------------------------------------------------
NAT Policy: nat-pol-1
Outside IP: 203.0.113.1 (vprn2)
Ports     : 1024-5119

-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[1/1/1:81] - sla:sla-prof-1
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
192.168.1.110
                00:0c:29:00:00:11    IPoE           DHCP         1          Y
2001:db8:101:1010::/64
                00:0c:29:00:00:11    IPoE           SLAAC        1          Y
192.168.1.100
                00:0c:29:00:00:14    IPoE           DHCP         1          Y
2001:db8:101:1010::/64
                00:0c:29:00:00:14    IPoE           SLAAC        1          Y
198.51.100.110
                00:0c:29:00:00:1e    IPoE           AAA          1          Y
192.168.1.254
```

```
                00:0c:29:00:00:1f    IPoE           AAA          1          Y
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Number of active subscribers : 1
-------------------------------------------------------------------------------
*A:BNG#
```

The following command shows the subscriber hierarchy for a single subscriber. This way it is apparent which host belongs to which IPoE session and on which SAP. The subscriber ID used is the MAC address of the BRG and is accompanied with the sub-profile used. The bridge ID is accompanied with the BRG profile. The NAT outside IP address is accompanied with the service number and the NAT policy.

```
*A:BNG# show service active-subscribers hierarchy subscriber "00:0c:29:00:00:10"

===============================================================================
Active Subscribers Hierarchy
===============================================================================
Hierarchy
-------------------------------------------------------------------------------
-- 00:0c:29:00:00:10 (sub-prof-1)
   |   brg-id: 00:0c:29:00:00:10 - brg-profile: brg-prof-1
   |   NAT Outside IP: 203.0.113.1 (vprn2) policy nat-pol-1
   |
   +-- sap:[1/1/1:81] - sla:sla-prof-1
       |
       |-- IPOE-session - mac:00:0c:29:00:00:11 - svc:1
       |   |
       |   |-- 192.168.1.110 - DHCP - L2Aware
       |   |
       |   +-- 2001:db8:101:1010::/64 - SLAAC
       |
       |-- IPOE-session - mac:00:0c:29:00:00:14 - svc:1
       |   |
       |   |-- 192.168.1.100 - DHCP - L2Aware
       |   |
       |   +-- 2001:db8:101:1010::/64 - SLAAC
       |
       |-- IPOE-session - mac:00:0c:29:00:00:1e - svc:1
       |   |
       |   +-- 198.51.100.110 - AAA
       |
       +-- IPOE-session - mac:00:0c:29:00:00:1f - svc:1
           |
           +-- 192.168.1.254 - AAA - L2Aware

===============================================================================
*A:BNG#
```

The following command shows the subscriber hosts on VPRN-1. All these hosts belong to the same subscriber, with subscriber ID 00:0c:29:00:00:10. The subscriber ID used is the MAC address of the BRG.

```
*A:BNG# show service id 1 subscriber-hosts

===============================================================
Subscriber Host table
===============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID Origin       Fwding State
---------------------------------------------------------------
[1/1/1:81]             00:0c:29:00:00:10
  192.168.1.100
    00:0c:29:00:00:14    N/A       DHCP         Fwding
[1/1/1:81]             00:0c:29:00:00:10
  192.168.1.110
    00:0c:29:00:00:11    N/A       DHCP         Fwding
[1/1/1:81]             00:0c:29:00:00:10
  192.168.1.254
    00:0c:29:00:00:1f    N/A       AAA          Fwding
[1/1/1:81]             00:0c:29:00:00:10
  198.51.100.110
    00:0c:29:00:00:1e    N/A       AAA          Fwding
[1/1/1:81]             00:0c:29:00:00:10
  2001:db8:101:1010::/64
    00:0c:29:00:00:11    N/A       IPoE-SLAAC   Fwding
[1/1/1:81]             00:0c:29:00:00:10
  2001:db8:101:1010::/64
    00:0c:29:00:00:14    N/A       IPoE-SLAAC   Fwding
---------------------------------------------------------------
Number of subscriber hosts : 6
===============================================================
*A:BNG#
```

The DHCP lease state table on VRPN-1 can then be shown with the following command. These addresses are taken from the Alc-Home-Aware-Pool as defined by the RADIUS server. The 192.168.1.110 address is a sticky address returned in the Alc-Reserved-Addresses attribute. Obviously, no entries appear for the static hosts.

```
*A:BNG# show service id 1 dhcp lease-state

===============================================================================
DHCP lease state table, service 1
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LeaseTime  Origin   Stdby
-------------------------------------------------------------------------------
192.168.1.100   00:0c:29:00:00:14 [1/1/1:81]          03h11m20s  Radius
192.168.1.110   00:0c:29:00:00:11 [1/1/1:81]          03h11m18s  Radius
-------------------------------------------------------------------------------
Number of lease states : 2
===============================================================================
*A:BNG#
```

Because a public static IP address is returned by the RADIUS server when authenticating the BRG, a ‘static’ BRG host was created, and the routing table is adjusted accordingly.

```
*A:BNG# show router 1 route-table

===============================================================================
Route Table (Service: 1)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age         Pref
      Next Hop[Interface Name]                                      Metric
-------------------------------------------------------------------------------
10.11.11.1/32                                   Local   Local     11h48m56s   0
       int-DHCP                                                        0
192.168.1.0/24                                  Local   Local     11h48m56s   0
       sub-int-1                                                       0
198.51.100.0/24                                 Local   Local     11h48m56s   0
       sub-int-1                                                       0
198.51.100.110/32                               Remote  Sub Mgmt  11h48m41s   0
       [grp-int-1-1]                                                   0
-------------------------------------------------------------------------------
No. of Routes: 4
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
       S = Sticky ECMP requested
===============================================================================
*A:BNG#
```

The following command shows the corresponding L2-aware hosts. Only three L2-aware hosts are created, using private IP addresses on the ‘inside’, and public address on the ‘outside’. Traffic for these hosts passes through the ISA. Traffic for the static public address (10.0.10.110) does not undergo NAT and does not pass through the ISA.

```
*A:BNG# show service nat l2-aware-hosts

===============================================================================
Layer-2-Aware NAT hosts
===============================================================================
Subscriber                  : 00:0c:29:00:00:10
Inside IP address           : 192.168.1.100
-------------------------------------------------------------------------------
Policy                      : nat-pol-1
Outside router              : 2
Outside IP address          : 203.0.113.1

Subscriber                  : 00:0c:29:00:00:10
Inside IP address           : 192.168.1.110
-------------------------------------------------------------------------------
Policy                      : nat-pol-1
Outside router              : 2
Outside IP address          : 203.0.113.1

Subscriber                  : 00:0c:29:00:00:10
Inside IP address           : 192.168.1.254
-------------------------------------------------------------------------------
Policy                      : nat-pol-1
Outside router              : 2
Outside IP address          : 203.0.113.1
```

```
-------------------------------------------------------------------------------
No. of hosts: 3
===============================================================================
*A:BNG#
```

The following command shows the corresponding IPoE sessions. There is one session per device/MAC-address, so there are four IPoE sessions for this subscriber.

```
*A:BNG# show service id 1 ipoe session

===============================================================================
IPoE sessions for svc-id 1
===============================================================================
Sap Id                          Mac Address         Up Time         MC-Stdby
    Subscriber-Id
        [CircuitID] | [RemoteID]
-------------------------------------------------------------------------------
[1/1/1:81]                      00:0c:29:00:00:11   0d 11:48:42
    00:0c:29:00:00:10
[1/1/1:81]                      00:0c:29:00:00:14   0d 11:48:40
    00:0c:29:00:00:10
[1/1/1:81]                      00:0c:29:00:00:1e   0d 11:48:42
    00:0c:29:00:00:10
[1/1/1:81]                      00:0c:29:00:00:1f   0d 11:48:42
    00:0c:29:00:00:10
-------------------------------------------------------------------------------
CID | RID displayed when included in session-key
Number of sessions : 4
===============================================================================
*A:BNG#
```

With this single subscriber connected, the following command shows the managed hosts that are created.

```
*A:BNG# show service id 1 managed-hosts type aaa

===============================================================================
Managed aaa hosts
===============================================================================
IP address                              MAC address
-------------------------------------------------------------------------------
192.168.1.254/32                        00:0c:29:00:00:1f
198.51.100.110/32                       00:0c:29:00:00:1e
-------------------------------------------------------------------------------
No. of Managed hosts: 2
===============================================================================
*A:BNG#
```

Conclusion

Using a BRG in the home network instead of a full-fledged L3 RGW offers network operators a view on the IP addresses and MAC addresses used by home devices, enabling them to provide per-device service offerings. Integrating BRGs in their networks can help the operators to optimize the revenue stream.

WiFi Aggregation and Offload — Basic Open SSID

This chapter provides information about WiFi Aggregation and Offload — Basic Open SSID.

Topics in this chapter include:

Applicability

This chapter is applicable to all 7750 SR platforms that function as Wireless Local Area Network gateways (WLAN-GWs). The configuration uses soft-Generic Routing Encapsulation (GRE) as the access technology, which requires one or more WLAN-IOMs (IOM3-XP and 2 x MS-ISAs), and was tested using SR OS release 12.0.R5.

Overview

WiFi Aggregation and Offload functionality for the 7750 SR is supported on SR OS 10.0.R3 or later. The functionality includes a RADIUS proxy server with a RADIUS proxy cache and support for soft-GRE tunnels.

WLAN-GW subscribers are implemented using Enhanced Subscriber Management (ESM) on the Control Processing Module (CPM), to benefit from the extensive ESM features available on the 7750 SR platform. Many different WiFi Offload configurations are possible, with the two most versatile configurations being open and secure Service Set Identifier (SSID).

This configuration should be used as a starting point for operators who need to offer an open SSID, where any client can connect to an Access Point (AP) and obtain an IP address without authentication. In most cases, operators want users to go through an authentication process before allowing full Internet access using the open SSID; therefore, this configuration also includes a web portal.

IP address assignment and Internet connectivity can be achieved using various methods in SR OS. In this configuration, a local DHCP server provides IP addresses to the User Equipment (UE) and routing to the Internet is performed using Global Routing Table (GRT) leaking.

Several considerations typically affect the choice of a WiFi Offload solution. Will access be free or paid? Will equipment be preconfigured, or will users bring their own WiFi device? When there is no pre-existing subscription, an open SSID is the most obvious solution. To provide a paywall or to have the user acknowledge certain terms of use due to legal reasons, a web portal may also be required.

When a web portal is implemented, users who connect to the open SSID which are not yet authenticated have all their web traffic redirected to the web portal landing page. This is performed using an http-redirect filter applied to the initial (limited) Service Level Agreement (SLA) profile assigned to the UE. Typically, the operating system of the UE will detect the presence of the web portal and automatically open the login page for the user. When the user logs in, the web portal sends a RADIUS Change of Authorization (CoA) request to the WLAN-GW, changing the SLA profile to one that does not contain an http-redirect filter.

Besides authentication, a major consideration is the method used to achieve Internet connectivity. Will the users require public addresses or are private addresses sufficient? In case few public IP addresses are available, private IP addresses can be assigned to UEs and the WLAN-GW can perform a Network Address Translation (NAT) function. If public routable IP addresses can be made available to all UEs, traffic from the UEs can be routed by the WLAN-GW to the Internet.

When a UE connects to an open SSID (as shown in Figure 1), typically the UE attempts to obtain an IP address using Dynamic Host Configuration Protocol (DHCP). The WLAN-GW can serve as a DHCP relay or proxy and may obtain the IP address from an external source, or use a local DHCP server function. A DHCP Discover or Request packet from a UE will trigger a form of authentication where the WLAN-GW requests information about the UE, such as SLA profile or DHCP local pool name. This authentication is separate from the web portal authentication and occurs immediately when a UE connects.

In summary:

- DHCP Discover triggering RADIUS authentication
- DHCP completes and UE has SLA profile with limited access
- UE logs into a web portal
- Successful login causes the portal to send a RADIUS CoA which assigns an SLA profile with full access

Figure 226 **Call Flow for Open SSID**

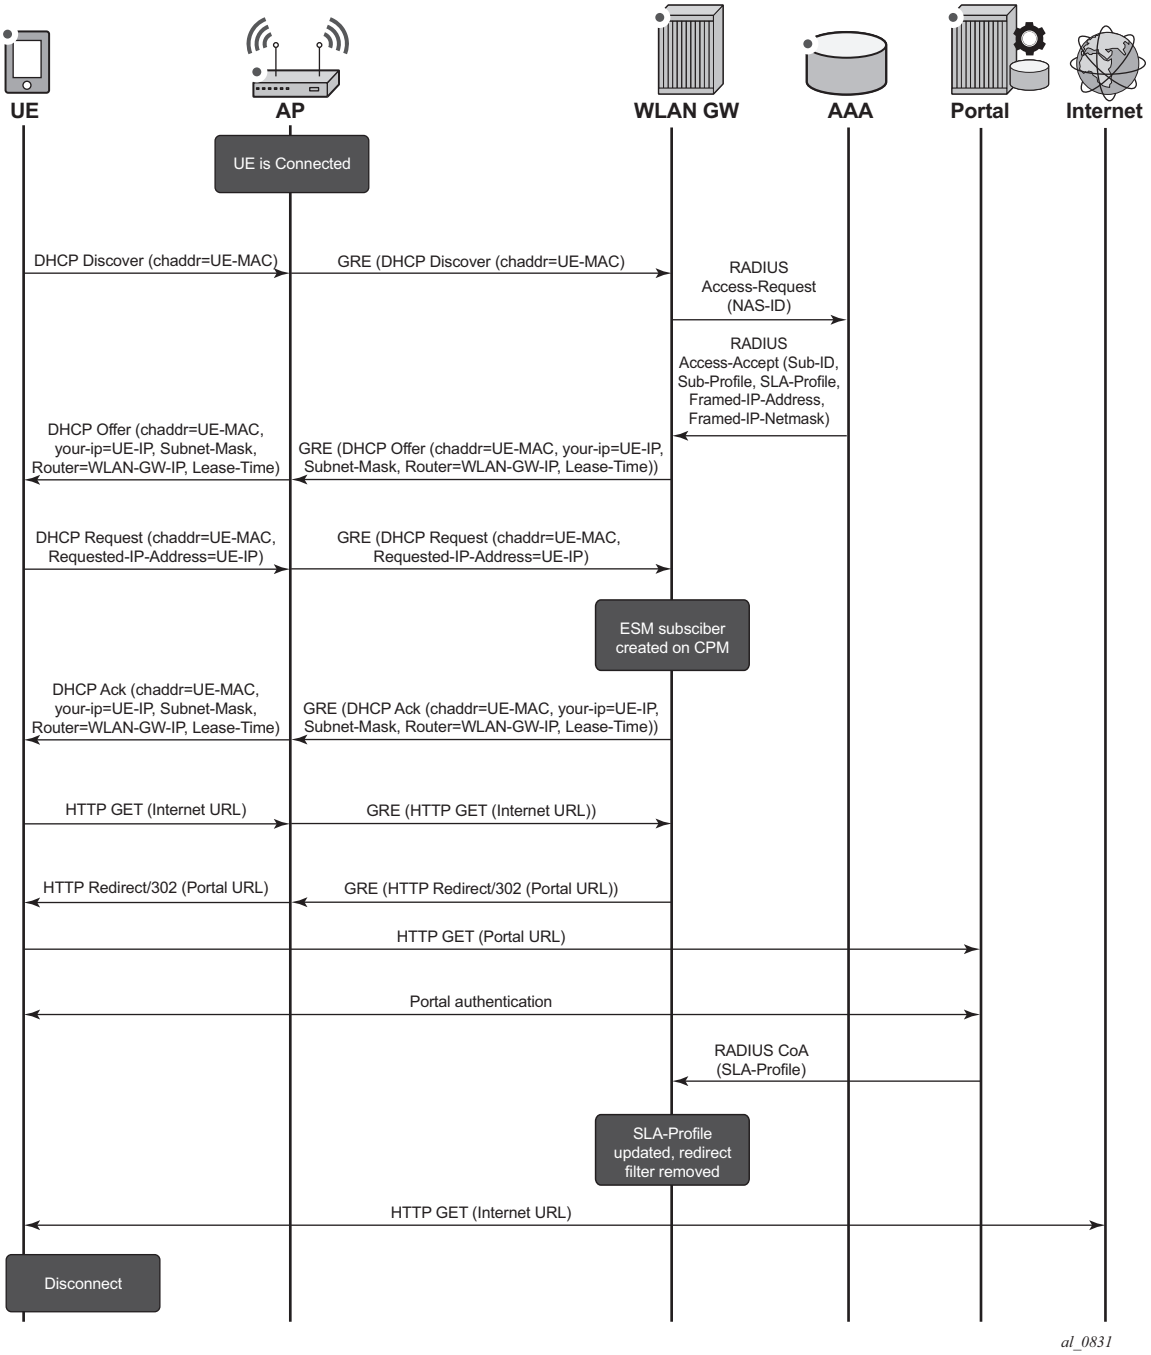

The SR OS is flexible in allowing the operator to separate the various WiFi Offload functions between different routing instances. All functions can be configured in the same routing instance, or as shown in Figure 2, the connectivity to the APs (and soft-GRE tunnels) can be provided in one Virtual Private Routed Network (VPRN), the users can be instantiated in another VPRN, and Authentication, Authorization and Accounting (AAA) access can be provided in yet another routing instance (in this case, the Base router). This clear separation of functions can enhance security; for example, by separating user traffic from authentication traffic.

Configuration

The WiFi offload scenario shown in Figure 2 has following characteristics:

- Open SSID with web portal authentication
- Local breakout to Internet using GRT leaking, routing through Base routing instance
- Same private IP address assigned to all UEs, with L2-aware NAT
- AP access in VPRN 1000
- UEs terminated in VPRN 2005
- Local DHCP server assigning public IP addresses

Figure 227 **WiFi Offload Scenario with Open SSID and Local DHCP Server**

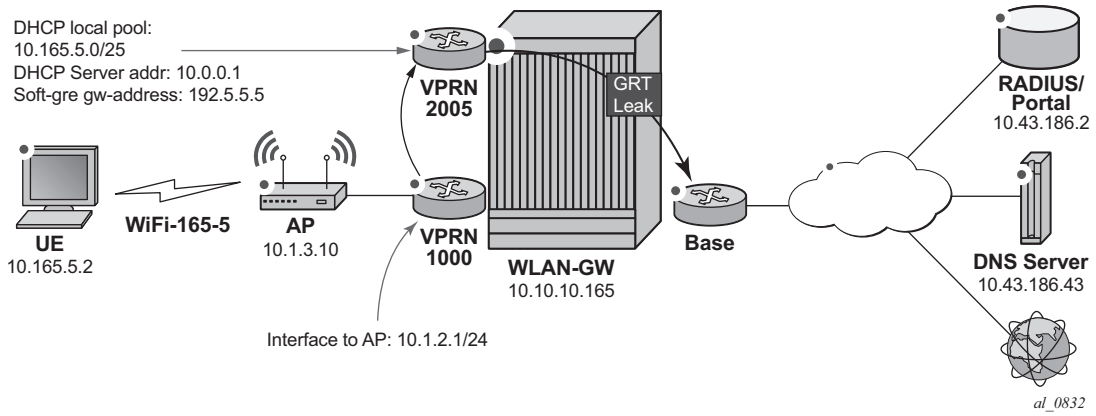

WLAN-GW

Note that the uplink interface, Interior Gateway Protocol (IGP), and system configuration is outside the scope of this document.

The following card and Media Dependent Adapter (MDA) configuration shows only the WLAN-Input/Output Module (IOM). An IOM3-XP-B containing two Multi-Service Integrated Service Adapter (MS-ISA) cards provides the WLAN-GW functionality. The MDA type for the ISA cards is isa-bb, the same type that is used for NAT.

```
*A:WLAN-GW# /configure card 2
*A:WLAN-GW>config>card# info
----------------------------------------------
        card-type iom3-xp-b
        mda 1
            mda-type isa-bb
            no shutdown
        exit
        mda 2
            mda-type isa-bb
            no shutdown
        exit
        no shutdown
----------------------------------------------
```

The following ISA configuration defines a wlan-gw-group referencing the IOM in slot two which hosts the two MS-ISA cards and providing the WLAN-GW functions.

```
A:WLAN-GW# /configure isa
A:WLAN-GW>config>isa# info
----------------------------------------------
        wlan-gw-group 1 create
            active-iom-limit 1
            iom 2
            no shutdown
        exit
----------------------------------------------
```

The following is a RADIUS server configuration, where the secret must match the secret configured on the external RADIUS server. The accept-coa option must be configured to allow the change of SLA profile by the web portal using a CoA request.

```
*A:WLAN-GW# /configure router radius-server
*A:WLAN-GW>config>router>radius-server# info
----------------------------------------------
            server "Server2" address 10.43.186.2 secret "zmLYVgt8UOLypJamceNSSHDWbZp
roq7Y" hash2 create
                accept-coa
            exit
----------------------------------------------
```

The following AAA configuration contains a RADIUS server policy used in the authentication policy. The source address must match the IP address configured for this client on the RADIUS server.

```
*A:WLAN-GW# /configure aaa
*A:WLAN-GW>config>aaa# info
----------------------------------------------
        radius-server-policy "RS_5" create
            servers
                router "Base"
                source-address 10.10.10.165
                server 1 name "Server2"
            exit
        exit
----------------------------------------------
```

The following policy configuration is used for exporting routes so that they are reachable by the public network. These policies are used for exporting UE routes in subsequent configuration sections.

```
*A:WLAN-GW# /configure router policy-options
*A:WLAN-GW>config>router>policy-options# info
----------------------------------------------
            prefix-list "WiFi-clients"
                prefix 10.165.0.0/16 longer
            exit
            policy-statement "WiFi-clients"
                entry 10
                    from
                        prefix-list "WiFi-clients"
                    exit
                    action accept
                    exit
                exit
            exit
----------------------------------------------
```

The uplink network configuration is outside the scope of this document. However, note that the IGP (here ISIS) must be aware of the UE addresses so that they are accessible from the Internet.

```
*A:WLAN-GW# /configure router isis
*A:WLAN-GW>config>router>isis# info
----------------------------------------------
        export "WiFi-clients"
----------------------------------------------
```

The following IP filter redirects all HTTP traffic to the web portal. The filter should also allow DNS and potentially other traffic, so the entry that allows TCP port 80 traffic to the web portal address must be placed before the entry that redirects all traffic to that portal; otherwise, there will be a redirect loop.

The HTTP redirect URL also includes a parameter that provides the MAC address of the UE to the web portal. In this configuration, either $MAC or $SUB can be used since both variables contain the MAC address of the UE. The web portal can reply with a CoA request specifying this particular UE MAC as the Subscriber ID after successful login. The URL also returns the IP address of the WLAN-GW to the web portal, so that the portal knows which WLAN-GW to send the CoA request to.

```
*A:WLAN-GW# /configure filter
*A:WLAN-GW>config>filter# info
----------------------------------------------
        ip-filter 2005 create
            default-action forward
            entry 70 create
                match protocol udp
                    dst-port eq 53
                exit
                action forward
            exit
            entry 80 create
                match protocol icmp
                exit
                action forward
            exit
            entry 90 create
                match protocol tcp
                    dst-ip 10.43.186.2/32
                    dst-port eq 80
                exit
                action forward
            exit
            entry 100 create
                match protocol tcp
                    dst-port eq 80
                exit
                action http-redirect "http://portal2.3ls.net/portal-no-
login.php?gw=10.10.10.165&mac=$SUB"
            exit
        exit
```

The following is a subscriber management configuration, with the RADIUS authentication policy used to authenticate DHCP requests, including the accept-authorization-change option to allow for SLA profile change after portal authentication. This DHCP authentication request also sends the NAS ID attribute that allows the RADIUS server to match on the configuration for this particular SSID. All UEs will be authenticated with their MAC address as user name, and **alcatel** as their password (any DHCP request will result in a successful authentication).

Two SLA profiles are required: profile SLAP_5_portal is initially used for each UE and refers to the portal redirect filter, while profile SLAP is applied using a CoA request after the user successfully authenticates on the web portal. A subscriber identity policy is also required.

```
configure subscriber-mgmt
        authentication-policy "WiFi-165-5-auth-policy" create
            password alcatel
            accept-authorization-change
            include-radius-attribute
                nas-identifier
            exit
            radius-server-policy "RS_5"
        exit
        sla-profile "SLAP" create
        exit
        sla-profile "SLAP_5_portal" create
            ingress
                ip-filter 2005
            exit
        exit
        sub-profile "SUBP" create
        exit
        sub-ident-policy "SIP" create
            sub-profile-map
                use-direct-map-as-default
            exit
            sla-profile-map
                use-direct-map-as-default
            exit
        exit
```

The following VPRN 1000 configuration contains the interface to the AP, and has GRT lookup with export-grt configured to allow APs to be managed from the Base routing instance.

```
*A:WLAN-GW# /configure service vprn 1000
*A:WLAN-GW>config>service>vprn# info
----------------------------------------------
            route-distinguisher 65400:1000
            interface "toAP3" create
                address 10.1.3.1/24
                sap 1/1/10 create
                exit
            exit
            grt-lookup
                enable-grt
                    static-route 0.0.0.0/0 grt
                exit
                export-grt "WiFi-APs"
            exit
            no shutdown
----------------------------------------------
```

VPRN 2005 is used for UE termination and contains:

- A local DHCP server with a single pool of addresses that are assigned to UEs.
- A loopback interface used by the DHCP server.
- A subscriber interface and group interface of type **wlangw** (called softgre prior to release 12.0).

- Subscriber parameters.
- The authentication policy, which will run each time a UE requests a DHCP address.
- The host-connectivity-verify function, which periodically checks the presence of UEs and quickly removes disconnected UEs even before their DHCP lease expires; the WLAN-GW has no other way of knowing when a UE has disconnected from the AP.
- The wlan-gw CLI-node (called soft-gre prior to release 12.0), including the wlan-gw GRE tunnel end-point address,and the routing instance where AP traffic is terminated, the ISA WLAN-GW group, and mobility parameters, which allow the UE state to be kept if the UE moves between two APs broadcasting the same SSID.
- GRT lookup with export-grt configured to allow UE traffic to be routed to the Internet.

```
*A:WLAN-GW# /configure service vprn 2005
*A:WLAN-GW>config>service>vprn# info
----------------------------------------------
            description "WiFi-165-5 Open SSID"
            dhcp
                local-dhcp-server "local_dhcp_2005" create
                    use-pool-from-client
                    pool "pool1" create
                        max-lease-time hrs 1
                        options
                            dns-server 10.43.186.43
                        exit
                        subnet 10.165.5.0/25 create
                            options
                                subnet-mask 255.255.255.128
                                default-router 10.165.5.1
                            exit
                            address-range 10.165.5.2 10.165.5.99
                        exit
                    exit
                    no shutdown
                exit
            exit
            route-distinguisher 65400:2005
            interface "dhcp-server" create
                address 10.0.0.1/24
                local-dhcp-server "local_dhcp_2005"
                loopback
            exit
            subscriber-interface "SI5" create
                address 10.165.5.1/24
                group-interface "GI5" wlangw create
                    sap-parameters
                        sub-sla-mgmt
                            def-sla-profile "SLAP_5_portal"
                            def-sub-profile "SUBP"
                            sub-ident-policy "SIP"
                        exit
                    exit
```

```
                    dhcp
                        option
                            action replace
                            circuit-id
                            no remote-id
                            vendor-specific-option
                                pool-name
                            exit
                        exit
                        server 10.0.0.1
                        trusted
                        lease-populate 10000
                        gi-address 10.165.5.1
                        no shutdown
                    exit
                    authentication-policy "WiFi-165-5-auth-policy"
                    host-connectivity-verify interval 5 action remove
                    wlan-gw
                        gw-address 192.5.5.5
                        mobility
                            trigger data iapp
                        exit
                        router 1000
                        wlan-gw-group 1
                        no shutdown
                    exit
                exit
            exit
            grt-lookup
                enable-grt
                exit
                export-grt "WiFi-clients"
            exit
            no shutdown
```

Freeradius

This simple default configuration section matches on any host. During the DHCP authentication phase, RADIUS returns the DHCP pool name **pool1** informing the WLAN-GW DHCP server which pool to assign the UE IP address from:

```
/etc/freeradius/users
DEFAULT         Auth-Type := Local, User-Password := "alcatel", user-name=~""
                        Alc-Subsc-ID-Str = "%{User-Name}",
                        Framed-Pool = "pool1",
```

In /etc/freeradius/clients.conf, the secret must match the secret configured in the WLAN-GW RADIUS server configuration.

```
client 10.10.10.165 {
        secret      = alcatel
        shortname   = WLAN-GW
}
```

The RADIUS CoA sent during successful portal login allows this UE full access, by applying SLA profile SLAP which does not have an http-redirect filter.

```
echo "ALC-Subsc-Id-Str='68:7f:74:8b:3d:d7',ALC-Subsc-Prof-Str='SUBP_5',ALC-SLA-Prof-
Str='.SLAP',Alc-Primary-Dns = '10.43.186.43'" | /usr/bin/radclient -x -r 1 -
t 2 '10.10.10.165' coa 'alcatel'
```

Access Points

At a minimum, the following must be configured on the Access Point:

- IP address 10.1.3.10/24
- Default route to 10.1.3.1
- Open SSID WiFi-165-5 mapped to VLAN 50
- Soft-GRE tunnel with destination 192.5.5.5, with VLAN 50 mapped to this tunnel

Show Commands

The following show commands reflect the status of the router after the UE has connected and obtained an IP address using DHCP.

The following output displays the UEs presently connected.

```
*A:WLAN-GW# show subscriber-mgmt wlan-gw ue
===============================================================================
User Equipments
===============================================================================
MAC address                 : 68:7f:74:8b:3d:d7
-------------------------------------------------------------------------------
VLAN Q-tag                  : 50
MPLS label                  : (Not Specified)
Tunnel router               : 1000
Tunnel remote IP address    : 10.1.3.10
Tunnel local IP address     : 192.5.5.5
Retail service              : N/A
SSID                        : (Not Specified)
Previous Access Point IP    : (Not Specified)
IMSI                        : (Not Specified)
Last move time              : 2014/09/22 10:47:58

-------------------------------------------------------------------------------
No. of UE: 1
===============================================================================
```

The DHCP lease information indicates that the address was assigned by the local DHCP server.

```
*A:WLAN-GW# show service id 2005 dhcp lease-state
===============================================================================
DHCP lease state table, service 2005
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LeaseTime  Origin   Stdby
-------------------------------------------------------------------------------
10.165.5.2      68:7f:74:8b:3d:d7 [2/1/nat-out-ip:20* 00h59m27s  DHCP
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

DHCP statistics can be displayed using following command.

```
*A:WLAN-GW# show service id 2005 dhcp statistics
===============================================================
DHCP Global Statistics, service 2005
===============================================================
Rx Packets                          : 2
Tx Packets                          : 2
Rx Malformed Packets                : 0
Rx Untrusted Packets                : 0
Client Packets Discarded            : 0
Client Packets Relayed              : 2
Client Packets Snooped              : 0
Client Packets Proxied (RADIUS)     : 0
Client Packets Proxied (User-Db)    : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Discarded            : 0
Server Packets Relayed              : 2
Server Packets Snooped              : 0
DHCP RELEASEs Spoofed               : 0
DHCP FORCERENEWs Spoofed            : 0
===============================================================
```

The route table for the routing instance where UEs are terminated shows an entry for the UE.

```
*A:WLAN-GW# show router 2005 route-table
===============================================================================
Route Table (Service: 2005)
===============================================================================
Dest Prefix[Flags]                              Type    Proto     Age        Pref
      Next Hop[Interface Name]                                     Metric
-------------------------------------------------------------------------------
10.0.0.0/24                                     Local   Local     00h32m12s  0
       dhcp-server                                                    0
10.165.5.0/24                                   Local   Local     00h27m20s  0
       SI5                                                            0
10.165.5.2/32                                   Remote  Sub Mgmt  00h00m33s  0
       [GI5]                                                          0
-------------------------------------------------------------------------------
No. of Routes: 3
Flags: n = Number of times nexthop is repeated
       B = BGP backup route available
       L = LFA nexthop available
```

```
       S = Sticky ECMP requested
===============================================================================
```

The active subscribers view shows the initial limited SLA profile SLAP_5_portal before the user has logged in to the portal.

```
*A:WLAN-GW>config>service>vprn# show service active-subscribers
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber 68:7f:74:8b:3d:d7 (SUBP)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[2/1/nat-out-ip:2049.3] - sla:SLAP_5_portal
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.165.5.2
                68:7f:74:8b:3d:d7 N/A       DHCP

-------------------------------------------------------------------------------
Number of active subscribers : 1
-------------------------------------------------------------------------------
```

The following output shows the active subscribers view after the user has logged in and the SLA profile has been updated with the unrestricted SLA profile SLAP.

```
*A:WLAN-GW# show service active-subscribers
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber 68:7f:74:8b:3d:d7 (SUBP_5)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[2/1/nat-out-ip:2049.3] - sla:SLAP
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.165.5.2
                68:7f:74:8b:3d:d7 N/A       DHCP

-------------------------------------------------------------------------------
Number of active subscribers : 1
-------------------------------------------------------------------------------
```

The following output shows the RADIUS statistics for the DHCP authentication.

```
*A:WLAN-GW# show aaa radius-server-policy "RS_5" statistics

===============================================================================
RADIUS server policy "RS_5" statistics
===============================================================================
```

```
Tx transaction requests                          : 1
Rx transaction responses                         : 1
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 "Server2" address 10.43.186.2 auth-port 1812 acct-port 1813
-------------------------------------------------------------------------------
Tx request packets                               : 1
Rx response packets                              : 1
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    :   7.24   7.24   7.24   7.24
Avg acct response delay (10 100 1K 10K) in ms    : n/a
Statistics last reset at                         : n/a

===============================================================================
```

The following output shows the CoA statistics after portal authentication.

```
*A:WLAN-GW# show subscriber-mgmt authentication coa-statistics
===============================================================================
Radius Notify Statistics    Change-Of-Authorization   Disconnect-Messages
===============================================================================
Requests Received             1                         0
Requests Accepted             1                         0
Requests Rejected             0                         0
Requests Dropped              0                         0
    No Auth Policy found      0                         0
    Invalid message           0                         0
    Out of resources          0                         0
    Authentication failure    0                         0
===============================================================================
```

Debug

The following is a complete debug of a UE connecting and logging in to the portal. Shortly after logging in, the UE disconnects from the SSID and the subscriber is removed by host-connectivity-verify.

The following debug configuration applies:

```
debug
    router "Base"
        radius
            packet-type authentication accounting coa
            detail-level medium
        exit
    exit
    router "2005"
        ip
            dhcp
                detail-level medium
                mode egr-ingr-and-dropped
            exit
        exit
        local-dhcp-server "local_dhcp_2005"
            detail-level medium
            mode dropped-only
        exit
    exit
    service
        id 2005
            host-connectivity-verify
                mac 68:7f:74:8b:3d:d7
            exit
        exit
    exit
exit
```

The WLAN-GW is notified of the UE after receiving the first DHCP packet.

```
1 2014/09/22 10:47:58.46 EDT MINOR: DEBUG #2001 vprn2005 PIP
"PIP: DHCP
instance 5 (2005), interface index 3 (GI5),
   received DHCP Boot Request on Interface GI5 (2/1/nat-out-ip:2049.3) Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 68:7f:74:8b:3d:d7    xid: 0x8c0fc642

   DHCP options:
   [53] Message type: Discover
   [61] Client id: (hex) 01 68 7f 74 8b 3d d7
   [50] Requested IP addr: 10.165.5.2
   [12] Host name: W81VM
   [60] Class id: MSFT 5.0
   [55] Param request list: len = 13
             1  Subnet mask
            15  Domain name
             3  Router
             6  Domain name server
            44  NETBIOS name server
            46  NETBIOS type
            47  NETBIOS scope
            31  Router discovery
            33  Static route
```

```
           121  Unknown option
           249  Unknown option
           252  Unknown option
            43  Vendor specific
   [255] End
"
```

DHCP triggers sending the RADIUS Access-Request.

```
2 2014/09/22 10:47:58.48 EDT MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 10.43.186.2:1812 id 1 len 79 vrid 1 pol RS_5
    USER NAME [1] 17 68:7f:74:8b:3d:d7
    PASSWORD [2] 16 IyDg9t17sGTbfR/6h0Bs1U
    NAS IP ADDRESS [4] 4 10.10.10.165
    NAS IDENTIFIER [32] 14 WLAN-GW
"
```

The UE authentication request is always accepted and the Access-Accept message contains the required subscriber management and IP parameters, in this case, at least the subscriber ID string as well as the pool name to be used by the local DHCP server.

```
3 2014/09/22 10:47:58.50 EDT MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 1 len 52 from 10.43.186.2:1812 vrid 1 pol RS_5
    VSA [26] 19 Alcatel(6527)
      SUBSC ID STR [11] 17 68:7f:74:8b:3d:d7
    FRAMED POOL [88] 5 pool1
"
```

The DHCP request is transmitted to the local DHCP server, which assigns the IP address to the UE.

```
4 2014/09/22 10:47:58.50 EDT MINOR: DEBUG #2001 vprn2005 PIP
"PIP: DHCP
instance 5 (2005),
   transmitted DHCP Boot Request to 10.0.0.1 Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 10.165.5.1
   chaddr: 68:7f:74:8b:3d:d7    xid: 0x8c0fc642

   DHCP options:
   [82] Relay agent information: len = 54
      [1] Circuit-id: WLAN-GW|2005|GI5|2/1/nat-out-ip:2049.3
      [9] Vendor-Specific info: len = 12
          Enterprise [6527] : len = 7
          [13] dhcpPool: pool1
   [53] Message type: Discover
   [61] Client id: (hex) 01 68 7f 74 8b 3d d7
   [50] Requested IP addr: 10.165.5.2
   [12] Host name: W81VM
   [60] Class id: MSFT 5.0
```

```
   [55] Param request list: len = 13
             1  Subnet mask
            15  Domain name
             3  Router
             6  Domain name server
            44  NETBIOS name server
            46  NETBIOS type
            47  NETBIOS scope
            31  Router discovery
            33  Static route
           121  Unknown option
           249  Unknown option
           252  Unknown option
            43  Vendor specific
   [255] End
"

5 2014/09/22 10:47:58.50 EDT MINOR: DEBUG #2001 vprn2005 PIP
"PIP: DHCP
instance 5 (2005),
   received DHCP Boot Reply on 10.0.0.1 Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.165.5.2
   siaddr: 10.0.0.1          giaddr: 10.165.5.1
   chaddr: 68:7f:74:8b:3d:d7    xid: 0x8c0fc642

   DHCP options:
   [82] Relay agent information: len = 54
      [1] Circuit-id: WLAN-GW|2005|GI5|2/1/nat-out-ip:2049.3
      [9] Vendor-Specific info: len = 12
          Enterprise [6527] : len = 7
          [13] dhcpPool: pool1
   [53] Message type: Offer
   [54] DHCP server addr: 10.0.0.1
   [51] Lease time: 3600
   [1] Subnet mask: 255.255.255.128
   [3] Router: 10.165.5.1
   [6] Domain name server: 10.43.186.43
   [12] Host name: W81VM
   [60] Class id: MSFT 5.0
   [255] End
"

6 2014/09/22 10:47:58.52 EDT MINOR: DEBUG #2001 vprn2005 PIP
"PIP: DHCP
instance 5 (2005), interface index 3 (GI5),
   transmitted DHCP Boot Reply to Interface GI5 (2/1/nat-out-ip:2049.3) Port 68

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.165.5.2
   siaddr: 10.0.0.1          giaddr: 10.165.5.1
   chaddr: 68:7f:74:8b:3d:d7    xid: 0x8c0fc642

   DHCP options:
   [53] Message type: Offer
   [54] DHCP server addr: 10.0.0.1
   [51] Lease time: 3600
   [1] Subnet mask: 255.255.255.128
```

```
   [3] Router: 10.165.5.1
   [6] Domain name server: 10.43.186.43
   [12] Host name: W81VM
   [60] Class id: MSFT 5.0
   [255] End
"

7 2014/09/22 10:47:58.69 EDT MINOR: DEBUG #2001 vprn2005 PIP
"PIP: DHCP
instance 5 (2005), interface index 3 (GI5),
   received DHCP Boot Request on Interface GI5 (2/1/nat-out-ip:2049.3) Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 68:7f:74:8b:3d:d7    xid: 0x8c0fc642

   DHCP options:
   [53] Message type: Request
   [61] Client id: (hex) 01 68 7f 74 8b 3d d7
   [50] Requested IP addr: 10.165.5.2
   [54] DHCP server addr: 10.0.0.1
   [12] Host name: W81VM
   [81] client FQDN: rcode1: 0, rcode2: 0, domain name = (hex) 00 57 38 31 56
   4d
   [60] Class id: MSFT 5.0
   [55] Param request list: len = 13
             1  Subnet mask
            15  Domain name
             3  Router
             6  Domain name server
            44  NETBIOS name server
            46  NETBIOS type
            47  NETBIOS scope
            31  Router discovery
            33  Static route
           121  Unknown option
           249  Unknown option
           252  Unknown option
            43  Vendor specific
   [255] End
"

8 2014/09/22 10:47:58.69 EDT MINOR: DEBUG #2001 vprn2005 PIP
"PIP: DHCP
instance 5 (2005),
   transmitted DHCP Boot Request to 10.0.0.1 Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 10.165.5.1
   chaddr: 68:7f:74:8b:3d:d7    xid: 0x8c0fc642

   DHCP options:
   [82] Relay agent information: len = 54
      [1] Circuit-id: WLAN-GW|2005|GI5|2/1/nat-out-ip:2049.3
      [9] Vendor-Specific info: len = 12
          Enterprise [6527] : len = 7
          [13] dhcpPool: pool1
```

```
   [53] Message type: Request
   [61] Client id: (hex) 01 68 7f 74 8b 3d d7
   [50] Requested IP addr: 10.165.5.2
   [54] DHCP server addr: 10.0.0.1
   [12] Host name: W81VM
   [81] client FQDN: rcode1: 0, rcode2: 0, domain name = (hex) 00 57 38 31 56
   4d
   [60] Class id: MSFT 5.0
   [55] Param request list: len = 13
             1  Subnet mask
            15  Domain name
             3  Router
             6  Domain name server
            44  NETBIOS name server
            46  NETBIOS type
            47  NETBIOS scope
            31  Router discovery
            33  Static route
           121  Unknown option
           249  Unknown option
           252  Unknown option
            43  Vendor specific
   [255] End
"

9 2014/09/22 10:47:58.69 EDT MINOR: DEBUG #2001 vprn2005 PIP
"PIP: DHCP
instance 5 (2005),
   received DHCP Boot Reply on 10.0.0.1 Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.165.5.2
   siaddr: 10.0.0.1          giaddr: 10.165.5.1
   chaddr: 68:7f:74:8b:3d:d7    xid: 0x8c0fc642

   DHCP options:
   [82] Relay agent information: len = 54
      [1] Circuit-id: WLAN-GW|2005|GI5|2/1/nat-out-ip:2049.3
      [9] Vendor-Specific info: len = 12
          Enterprise [6527] : len = 7
          [13] dhcpPool: pool1
   [53] Message type: Ack
   [54] DHCP server addr: 10.0.0.1
   [51] Lease time: 3600
   [1] Subnet mask: 255.255.255.128
   [3] Router: 10.165.5.1
   [6] Domain name server: 10.43.186.43
   [12] Host name: W81VM
   [81] client FQDN: rcode1: 0, rcode2: 0, domain name = (hex) 00 57 38 31 56
   4d
   [60] Class id: MSFT 5.0
   [255] End
"

10 2014/09/22 10:47:58.69 EDT MINOR: DEBUG #2001 vprn2005 PIP
"PIP: DHCP
instance 5 (2005), interface index 3 (GI5),
   transmitted DHCP Boot Reply to Interface GI5 (2/1/nat-out-ip:2049.3) Port 68
```

```
   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 10.165.5.2
   siaddr: 10.0.0.1          giaddr: 10.165.5.1
   chaddr: 68:7f:74:8b:3d:d7    xid: 0x8c0fc642

   DHCP options:
   [53] Message type: Ack
   [54] DHCP server addr: 10.0.0.1
   [51] Lease time: 3600
   [1] Subnet mask: 255.255.255.128
   [3] Router: 10.165.5.1
   [6] Domain name server: 10.43.186.43
   [12] Host name: W81VM
   [81] client FQDN: rcode1: 0, rcode2: 0, domain name = (hex) 00 57 38 31 56
   4d
   [60] Class id: MSFT 5.0
   [255] End
"
```

At this point in the configuration, the UE has network connectivity but all HTTP traffic is redirected to the web portal, as configured in the IP filter included in the initial SLA profile.

After web portal authentication, the WLAN-GW receives a RADIUS CoA for this subscriber, which includes the new unrestricted SLA profile SLAP.

```
11 2014/09/22 10:48:12.54 EDT MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Change of Authorization(43) id 162 len 71 from 10.43.186.2:55255 vrid 1
    VSA [26] 19 Alcatel(6527)
      SUBSC ID STR [11] 17 68:7f:74:8b:3d:d7
    VSA [26] 8 Alcatel(6527)
      SUBSC PROF STR [12] 6 SUBP_5
    VSA [26] 6 Alcatel(6527)
      SLA PROF STR [13] 4 SLAP
"

12 2014/09/22 10:48:12.54 EDT MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Change of Authorization Ack(44) 10.43.186.2:55255 id 162 len 20 vrid 1
"
```

The host has accessed a few web sites, then disconnected from the SSID, which is not known by the WLAN-GW. After 5 minutes of inactivity, host-connectivity-verify removes the subscriber and the DHCP lease is cleared.

```
13 2014/09/22 10:48:58.90 EDT MINOR: DEBUG #2001 vprn2005 SHCV
"SHCV: Periodic Check
   2/1/nat-out-ip:2049.3
   DHCP lease state 10.165.5.2 68:7f:74:8b:3d:d7"

14 2014/09/22 10:49:08.90 EDT MINOR: DEBUG #2001 vprn2005 SHCV
"SHCV: Periodic Check
   2/1/nat-out-ip:2049.3
   DHCP lease state 10.165.5.2 68:7f:74:8b:3d:d7"
```

```
15 2014/09/22 10:49:18.90 EDT MINOR: DEBUG #2001 vprn2005 SHCV
"SHCV: Periodic Check
   2/1/nat-out-ip:2049.3
   DHCP lease state 10.165.5.2 68:7f:74:8b:3d:d7"

16 2014/09/22 10:49:28.90 EDT MINOR: DEBUG #2001 vprn2005 SHCV
"SHCV: Connectivity Lost
   2/1/nat-out-ip:2049.3
   DHCP lease state 10.165.5.2 68:7f:74:8b:3d:d7"

17 2014/09/22 10:49:30.00 EDT MINOR: DEBUG #2001 vprn2005 PIP
"PIP: DHCP
instance 5 (2005),
   transmitted DHCP Boot Request to 10.0.0.1 Port 68

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 10.165.5.2        yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 68:7f:74:8b:3d:d7    xid: 0x0

   DHCP options:
   [53] Message type: Release
   [54] DHCP server addr: 10.0.0.1
   [255] End
"
```

Conclusion

The 7750 SR WLAN-GW can support many WiFi Offload architectures, including open SSID with portal authentication. WiFi Offload functions such as terminating GRE tunnels or subscribers can be performed in separate routing instances, if required. IP addresses can be assigned from an external or local source and routing can be performed using NAT, by connecting the UE routing instance directly to the Internet, or by leaking routes to other routing instances. Using http-redirect, a web portal can be used to allow users to log in to a paid service or to accept the terms of service for a free WiFi service. Several show commands and debug options are available to help the operator monitor and troubleshoot the solution.

WiFi Aggregation and Offload — Basic Secure SSID with Distributed RADIUS Proxy

This chapter provides information about WiFi Aggregation and Offload — Basic Secure SSID with Distributed RADIUS Proxy.

Topics in this chapter include:

Applicability

The configuration example in this chapter is applicable to all 7750 SR platforms that function as WLAN gateway (WLAN-GW). This configuration makes use of soft-GRE as the access technology which requires one or more WLAN-IOMs (IOM3-XP and two MS-ISAs) and was tested using SR OS release 13.0.R3.

Summary

WiFi Aggregation and Offload functionality for the 7750 SR has been supported in SR OS 10.0.R3, and later. This includes a RADIUS proxy server and support for soft-GRE tunnels.

WLAN-GW subscribers can be implemented using Enhanced Subscriber Management (ESM) on the CPM in order to benefit from the extensive ESM features available on the 7750 SR platform. Many different WiFi Offload configurations are possible, with two versatile categories being open and secure SSID.

Starting from SR OS release 12.0.R4, distributed RADIUS-proxy functionality (DRP) has been added to the MS-ISA. This feature allows running a high-performance proxy over multiple MS-ISA cards instead of being limited to a single CPM, greatly increasing scalability.

This chapter can be used as a starting point for operators who wish to configure a secure SSID scenario using DRP and ESM. In a secure SSID scenario, the Access Point (AP) uses 802.1x and Extensible Authentication Protocol (EAP) to authenticate the UE. The EAP method used is transparent to the WLAN-GW. In this chapter, PEAP/EAP-MSCHAPv2 is used to associate with the SSID by entering a user name and password, but other methods such as EAP-Subscriber Identity Module (EAP-SIM) can also be used without any configuration change on the WLAN-GW.

IP address assignment and Internet connectivity can be achieved by using various methods on the 7750. In this scenario, the RADIUS server provides the IP addresses to the User Equipment (UEs). The same IP private address is assigned to every UE and L2-aware Network Address Translation (NAT) is used to provide a public IP address on the Internet.

Overview

For a WiFi Offload solution where the service provider has a record of their users, that is, where users have login accounts, the provider may consider offering a secure SSID as a more convenient and secure alternative to an open SSID with a web portal. In a secure SSID scenario, all user traffic is encrypted between the UE and the AP, and UEs are only granted access if they authenticate successfully. This makes attacks more difficult and blocks non-paying users who only connect to test if they can get free access.

Authentication in this case requires a centralized Authentication, Authorization and Accounting (AAA) which keeps track of the user accounts. The user is granted full access immediately after connecting to the secure SSID. One drawback is that WiFi clients may need some configuring by the user before they are able to connect to the SSID using the correct EAP method. In the case of EAP-SIM, users do not need to know their user name and password because authentication is done based on credentials contained in the SIM card, but the SSID configuration may need to be preloaded by the operator on the mobile device or provided to the user ahead of time. For other EAP methods such as PEAP/EAP-MSCHAPv2, users need to supply the correct user name and password, without the help of a portal or any instructions to guide them.

An operator offering Internet access to a large number of users while only a limited number of public IP addresses are available will likely use Network Address Translation (NAT) in order to conserve public IP addresses. NAT typically maps a few public IP addresses and ports to a large number of inside (private) IP addresses and ports. The 7750 SR WLAN-GW supports several NAT configurations including L2-aware NAT, where the MAC address of the UE is also used when creating the mapping between the inside IP/port and the outside IP/port. Therefore with L2-aware

NAT, the same private IP address can be assigned to all UEs because the unique MAC address for each UE allows the WLAN-GW to distinguish between each UE. This greatly simplifies IP address assignment; the RADIUS server can assign the same private IP address to all UEs and there is no DHCP server required. Using a RADIUS server for IP address assignment means DHCP proxy needs to be configured on the WLAN-GW.

The 7750 SR platform is flexible in terms of allowing the operator to separate the various WiFi Offload functions between different routing instances. All functions can be configured in the same routing instance, or as shown in the following configuration, the connectivity to the APs (and soft-GRE tunnels) can be provided in one Virtual Private Routed Network (VPRN), the users can be instantiated in another VPRN, and AAA access can be provided in yet another routing instance (in this example the Base router). This provides a clear separation of functions and can enhance security, by separating user traffic from authentication and management traffic.

Configuration

The WiFi offload scenario with SSID and L2-aware NAT shown in Figure 228 has following characteristics:

- Secure SSID with EAP authentication
- Local breakout to Internet, routing through the Base routing instance
- Same private IP address assigned to all UEs by RADIUS
- L2-aware NAT
- AP access in VPRN 1000
- UEs terminated in VPRN 2001

Figure 228 **WiFi Offload Scenario with Secure SSID and L2-Aware NAT**

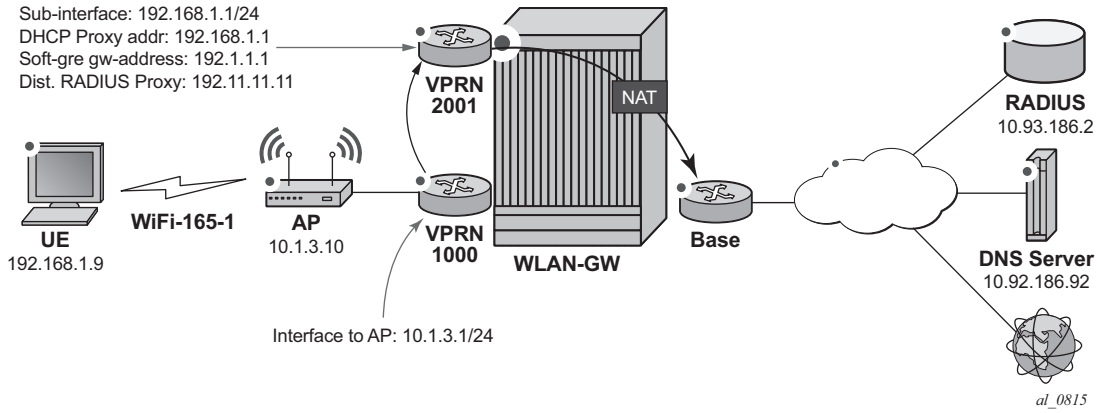

WLAN-GW

Note that configuring the uplink interface, Interior Gateway Protocol (IGP) and system configuration is outside the scope of this chapter and only partial configuration is provided.

Card and Media Dependent Adapter (MDA) configuration showing only the WLAN-IOM. An IOM3-XP-B containing two MS-ISA cards provides the WLAN-GW functionality. The MDA type for the ISA cards is isa-bb, the same type that is used for NAT.

```
*A:WLAN-GW# /configure card 2
*A:WLAN-GW>config>card# info
----------------------------------------------
        card-type iom3-xp-b
        mda 1
            mda-type isa-bb
            no shutdown
        exit
        mda 2
            mda-type isa-bb
            no shutdown
        exit
        no shutdown
----------------------------------------------
```

The ISA configuration contains a wlan-gw-group referencing the IOM in slot 2, which hosts the two MS-ISA cards providing the WLAN-GW functions.

```
A:WLAN-GW# /configure isa
A:WLAN-GW>config>isa# info
----------------------------------------------
        wlan-gw-group 1 create
```

```
            active-iom-limit 1
            iom 2
            no shutdown
        exit
-----------------------------------------------
```

The AAA configuration contains an ISA RADIUS policy used for authentication requests. The source address range configures the source address of the first MS-ISA in the wlan-gw-group. The second MS-ISA card gets the next consecutive IP address and so on. All the IP addresses assigned this way to MS-ISA cards must be configured as clients on the RADIUS server. The secret here must match the secret configured on the RADIUS server.

```
A:WLAN-GW# /configure aaa
A:WLAN-GW>config>aaa# info
----------------------------------------------
        isa-radius-policy "IRS_1" create
            servers
                router "Base"
                source-address-range 10.10.165.1
                server 1 create
                    authentication
                    ip-address 10.93.186.2
                    secret "7USmr6f7JkxaGnDDq1uqwEAJKGbhZr5i" hash2
                    no shutdown
                exit
            exit
        exit
-----------------------------------------------
```

The following policy shows the two routes that must be exported for this scenario to work: UE NAT outside routes (NAT is configured in the next step), to make UEs reachable on the Internet, and MS-ISA RADIUS source address routes, in order for the MS-ISAs to be reachable from the RADIUS server. This policy should be used for export in the IGP configuration (not shown).

```
A:WLAN-GW# /configure router policy-options
A:WLAN-GW>config>router>policy-options# info
----------------------------------------------
            prefix-list "WiFi"
                prefix 10.10.165.0/24 longer
                prefix 10.165.0.0/16 longer
            exit
            policy-statement "WiFi"
                entry 10
                    from
                        prefix-list "WiFi"
                    exit
                    action accept
                    exit
                exit
            exit
-----------------------------------------------
```

The following configures L2-aware NAT by creating an outside NAT pool with a public IP address range. The private inside address used by the UE will be mapped to an outside IP address routable on the Internet. NAT port mapping parameters can be set in this configuration, controlling how many outside ports can be used by each UE. Details of NAT configuration are outside the scope of this document.

```
A:WLAN-GW# /configure router nat
A:WLAN-GW>config>router>nat# info
----------------------------------------------
            outside
                pool "WiFi-165-1" nat-group 1 type l2-aware create
                    address-range 10.165.1.0 10.165.1.255 create
                    exit
                    no shutdown
                exit
            exit
----------------------------------------------
```

The following configures a NAT policy under services, linking this policy with the outside NAT pool. When the NAT policy is invoked for a subscriber, this associates the subscriber with the correct outside pool.

```
A:WLAN-GW# /configure service nat
A:WLAN-GW>config>service>nat# info
----------------------------------------------
            nat-policy "WiFi-165-1" create
                pool "WiFi-165-1" router Base
            exit
----------------------------------------------
```

The following subscriber management configuration includes an SLA profile, a subscriber identity policy, and the subscriber profile that makes use of the previously defined NAT policy. This allows subscriber traffic to be forwarded to the Internet through the Base routing instance where the outside NAT pool exists.

A dummy authentication-policy is required for the CPM to handle the DHCP Discover messages forwarded by the MS-ISA cards.

```
A:WLAN-GW# /configure subscriber-mgmt
A:WLAN-GW>config>subscr-mgmt# info
----------------------------------------------
        authentication-policy "dummy-auth-policy" create
        exit
        sla-profile "SLAP_1" create
        exit
        sub-profile "SUBP_1" create
            nat-policy "WiFi-165-1"
        exit
        sub-ident-policy "SIP" create
            sub-profile-map
                use-direct-map-as-default
            exit
            sla-profile-map
                use-direct-map-as-default
```

```
            exit
        exit
-----------------------------------------------
```

VPRN 1000 contains the interface to the AP as well as the distributed RADIUS proxy server RP_1. The RADIUS proxy wlan-gw address configures a special NAT route in VPRN 1000 that forwards RADIUS packets from the AP to the correct MS-ISA. That address is known to the AP as the RADIUS server address it uses for EAP authentication for this SSID. The secret configured here has to match the RADIUS secret configured on the AP.

The RADIUS proxy is configured to create cache entries based on attribute 31 in RADIUS access-request packets (Calling-Station-ID), which contains the MAC address of the UE. These cache entries will be stored temporarily and used to authenticate DHCP packets from the UE. The **track-accounting start** parameter allows mobility to be triggered for a UE upon receiving an accounting-start message. The UE's associated tunnel will be moved to the IP address indicated by the NAS-IP-Address. The **track-accounting stop** parameter allows the UE session to be terminated immediately when the AP sends a RADIUS accounting-stop for the UE, when this UE disconnects from the SSID.

The default-authentication-server-policy links the RADIUS proxy with the isa-radius-policy that authenticates the UEs. If accounting is required, the accounting policy can be specified in this configuration and can be the same as or different from the isa-radius-policy. The send-accounting-response option makes the WLAN-GW acknowledge (and then discard) the RADIUS accounting messages from the AP, instead of proxying the accounting messages to the external RADIUS server.

```
A:WLAN-GW# /configure service vprn 1000
A:WLAN-GW>config>service>vprn# info
----------------------------------------------
            route-distinguisher 65400:1000
            interface "toAP3" create
                address 10.1.3.1/24
                sap 1/1/10 create
                exit
            exit
            radius-proxy
                server "DRP_1" purpose accounting authentication wlan-gw-
group 1 create
                    cache
                        key packet-type request attribute-type 31
                        track-accounting start stop
                        no shutdown
                    exit
                    default-authentication-server-policy "IRS_1"
                    secret "nUeorYjgFZtuAqIwoUOLODFxF43rhSf/" hash2
                    send-accounting-response
                    wlan-gw
                        address 192.11.11.11
                    exit
                    no shutdown
```

```
                exit
            exit
            no shutdown
----------------------------------------------
```

VPRN 2001 is used for UE termination and contains:

- A subscriber interface and group interface of type **wlangw** (**softgre** prior to release 12.0).
- Default subscriber parameters assigned to every UE.
- DHCP proxy, which allows the RADIUS-assigned IP address parameters stored in the DRP cache during authentication to be passed to the UE.
- A dummy authentication policy which allows the CPM to handle the DHCP Discover passed on by the MS-ISA.
- The **wlan-gw** node (**soft-gre** prior to release 12.0), which includes:
  - The gw-address that is the end-point of the GRE tunnel
  - The routing instance where AP traffic is terminated
  - The ISA wlan-gw-group, which associates this WLAN-GW configuration with a set of IOMs
  - Mobility parameters, which allow the UE state to be kept if the UE moves between two APs broadcasting the same SSID
  - The authenticate-on-dhcp option required for the CPM to instantiate ESM UEs when using DRP
  - The L2-aware address/subnet used for L2-aware NAT. This address matches the default gateway assigned to the UEs.

```
A:WLAN-GW# /configure service vprn 2001
A:WLAN-GW>config>service>vprn# info
----------------------------------------------
            description "WiFi-165-1 Secure SSID"
            route-distinguisher 65400:2001
            subscriber-interface "SI1" create
                address 192.168.1.1/24 populate-host-routes
                group-interface "GI1" wlangw create
                    sap-parameters
                        sub-sla-mgmt
                            def-sla-profile "SLAP_1"
                            def-sub-profile "SUBP_1"
                            sub-ident-policy "SIP"
                        exit
                    exit
                    dhcp
                        proxy-server
                            emulated-server 192.168.1.1
                            no shutdown
                        exit
                        lease-populate 10000
                        gi-address 192.168.1.1
```

```
                        no shutdown
                    exit
                    authentication-policy "dummy-auth-policy"
                    wlan-gw
                        gw-address 192.1.1.1
                        mobility
                            trigger data iapp
                        exit
                        router 1000
                        wlan-gw-group 1
                        vlan-tag-ranges
                            range default
                                authenticate-on-dhcp
                            exit
                        exit
                        no shutdown
                    exit
                exit
            exit
            nat
                inside
                    l2-aware
                        address 192.168.1.1/24
                    exit
                exit
            exit
            no shutdown
-----------------------------------------------
```

Freeradius

This part of the user's configuration file matches on the user name entered by the UE while connecting to this secure SSID. If the password entered is correct, RADIUS returns the IP addressing parameters configured as follows. The same IP address 192.168.1.9 is assigned to every user on this SSID, but the L2-aware NAT on the WLAN-GW can distinguish between all the UEs based on their L2 MAC address.

```
/etc/freeradius/users:
"user1"                 User-Password := "alcatel"
                        Alc-Subsc-ID-Str = "%{User-Name}",
                        Alc-Default-Router = 192.168.1.1,
                        Alc-Primary-Dns = 10.43.186.43,
                        Framed-IP-Address = 192.168.1.9,
                        Framed-IP-Netmask = 255.255.255.0,
```

In /etc/freeradius/clients.conf the secret matches the one configured in the WLAN-GW isa-radius-policy configuration. Since there are only two MS-ISA cards in the wlan-group used in this example, two clients are configured.

```
client 10.10.165.1 {
        secret      = alcatel
        shortname   = WLAN-GW-ISA1
}
```

```
client 10.10.165.2 {
        secret      = alcatel
        shortname   = WLAN-GW-ISA2
}
```

Access Points

The following must be configured on the Access Point as a minimum:

- IP address 10.1.3.10/24
- Default route to 10.1.3.1
- Secure SSID WiFi-165-1 mapped to VLAN 10, using WPA2 with EAP/802.1x authenticating against RADIUS server 192.11.11.11, with RADIUS accounting enabled
- Soft-GRE tunnel with destination 192.1.1.1, with VLAN 10 mapped to this tunnel

Show Commands

The following show commands reflect the status of the WLAN-GW after the UE has connected and obtained and IP address using DHCP.

The following output displays the connected UEs:

```
A:WLAN-GW# show subscriber-mgmt wlan-gw ue
===============================================================================
User Equipments
===============================================================================
MAC address                 : 68:7f:74:8b:3d:d7
-------------------------------------------------------------------------------
VLAN Q-tag                  : 10
MPLS label                  : (Not Specified)
Tunnel router               : 1000
Tunnel remote IP address    : 10.1.3.10
Tunnel local IP address     : 192.1.1.1
Retail service              : N/A
SSID                        : "WiFi-165-1"
Previous Access Point IP    : (Not Specified)
IMSI                        : (Not Specified)
Subscriber host service     : 2001
Subscriber host SAP         : 2/1/nat-out-ip:2049.1
Last move time              : 2015/09/15 16:20:01

-------------------------------------------------------------------------------
No. of UE: 1
===============================================================================

A:WLAN-GW# tools dump wlan-gw ue
```

```
===============================================================================
Matched 1 session on Slot #2 MDA #1
===============================================================================
UE-Mac          : 68:7f:74:8b:3d:d7     UE-vlan         : 10
UE IP Addr      : N/A                   UE Timeout      : N/A
DHCPv6 Timeout  : N/A                   SLAAC Timeout   : N/A
DHCPv6 IA-NA ID : N/A                   RA Timeout      : N/A
DHCPv6 Addr     : N/A
SLAAC Prefix    : N/A
Description     : ESM-user
Auth/CoA-time   : 09/16/2015 10:47:38   Retail Service  : N/A
Tunnel MDA      : 2/2                   Tunnel Router   : 1000
MPLS label      : N/A                   Shaper          : 1
Tunnel Src IP   : 10.1.3.10             Tunnel Dst IP   : 192.1.1.1
Tunnel L2 Svc   : N/A                   Tunnel L2 Vlan  : N/A
Tunnel Type     : GRE
Anchor SAP      : 2/1/nat-out-ip:2049.2
AP-Mac          : 00:0d:67:39:0b:65     AP-RSSI         : Unknown
AP-SSID         : "WiFi-165-1"
Last-forward    : 09/16/2015 15:59:26   Last-move       : 09/16/2015 10:47:38
Session Timeout : None                  Idle Timeout    : N/A
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       : N/A                   Rx Octets       : N/A
Tx Frames       : N/A                   Tx Octets       : N/A
-------------------------------------------------------------------------------
===============================================================================
No sessions on Slot #2 MDA #2 match the query
```

The DHCP lease information indicates that the address was assigned by RADIUS.

```
A:WLAN-GW# show service id 2005 dhcp lease-state
===============================================================================
DHCP lease state table, service 2001
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LeaseTime  Origin   Stdby
-------------------------------------------------------------------------------
192.168.1.9     68:7f:74:8b:3d:d7 [2/1/nat-out-ip:20* 06d23h59m  Radius
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

When troubleshooting DHCP issues, displaying DHCP statistics is useful.

```
A:WLAN-GW# show service id 2005 dhcp statistics
===============================================================
DHCP Global Statistics, service 2001
===============================================================
```

```
Rx Packets                          : 2
Tx Packets                          : 2
Rx Malformed Packets                : 0
Rx Untrusted Packets                : 0
Client Packets Discarded            : 0
Client Packets Relayed              : 0
Client Packets Snooped              : 0
Client Packets Proxied (RADIUS)     : 2
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
===============================================================================
```

The following output lists the active subscribers, showing each UE SLA profile, MAC address and IP address.

```
A:WLAN-GW# show service active-subscribers
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber DUACBU2ZLE (SUBP_1)
-------------------------------------------------------------------------------
NAT Policy: WiFi-165-1
Outside IP: 10.165.1.0
Ports     : 1024-65535
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[2/1/nat-out-ip:2049.1] - sla:SLAP_1
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
192.168.1.9
                68:7f:74:8b:3d:d7 N/A       DHCP

-------------------------------------------------------------------------------
Number of active subscribers : 1
-------------------------------------------------------------------------------
```

The following output displays distributed RADIUS proxy server statistics after the UE has authenticated, showing all the EAP messages exchanged between the AP and RADIUS proxy:

```
A:WLAN-GW# show router 1000 radius-proxy-server "DRP_1" statistics
===============================================================================
ISA RADIUS Proxy server statistics for "DRP_1"
===============================================================================
Group 1 member 1
-------------------------------------------------------------------------------
Rx packet                                                  : 12
```

```
Rx Access-Request                                    : 11
Rx Accounting-Request                                : 1
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                                  : 10
Tx Accounting-Response                               : 1
Tx dropped                                           : 0
  Server timeout                                     : 0
  Invalid response Authenticator                     : 0
  Invalid Message-Authenticator                      : 0
  Invalid attribute encoding                         : 0
  RADIUS server send failure                         : 0

Group 1 member 2
-------------------------------------------------------------------------------
Rx packet                                            : 0
Rx Access-Request                                    : 0
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                                     : 0
Tx Access-Reject                                     : 0
Tx Access-Challenge                                  : 0
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 output shows the ISA RADIUS policy statistics after the UE has connected, showing the transactions between the WLAN-GW and the RADIUS server.

```
A:WLAN-GW# show aaa isa-radius-policy "IRS_1"
===============================================================================
ISA RADIUS policy "IRS_1"
===============================================================================
Description                 : (Not Specified)
Include attributes acct     : N/A
Include attributes auth     : nas-ip-address
User name format            : mac
User name MAC format        : alu
NAS-IP-Address              : system-ip
-------------------------------------------------------------------------------
RADIUS server settings
-------------------------------------------------------------------------------
Router                      : "Base"
Source address start        : 10.10.165.1
Source address end          : 10.10.165.2
Access algorithm            : direct
Retry                       : 3
Timeout (s)                 : 5
Last management change      : 09/15/2015 15:05:02
===============================================================================
===============================================================================
Servers for "IRS_1"
===============================================================================
Index Address         Acct-port Auth-port CoA-port
-------------------------------------------------------------------------------
1     10.93.186.2     0         1812      0
===============================================================================
===============================================================================
Status for ISA RADIUS server policy "IRS_1"
===============================================================================
Server 1, group 1, member 1
-------------------------------------------------------------------------------
Purposes Up                                 : authentication
Source IP address                           : 10.10.165.1
Acct Tx Requests                            : 0
Acct Tx Retries                             : 0
Acct Tx Timeouts                            : 0
Acct Rx Replies                             : 0
Auth Tx Requests                            : 11
Auth Tx Retries                             : 0
Auth Tx Timeouts                            : 0
Auth Rx Replies                             : 11
CoA Rx Requests                             : 0

Server 1, group 1, member 2
-------------------------------------------------------------------------------
Purposes Up                                 : (None)
Source IP address                           : 10.10.165.2
Acct Tx Requests                            : 0
Acct Tx Retries                             : 0
Acct Tx Timeouts                            : 0
Acct Rx Replies                             : 0
Auth Tx Requests                            : 0
```

```
Auth Tx Retries                                  : 0
Auth Tx Timeouts                                 : 0
Auth Rx Replies                                  : 0
CoA Rx Requests                                  : 0
===============================================================================
```

Call Flow

Figure 229 shows the call flow for a secure SSID with DRP. The main steps are:

- EAP authentication using DRP on MS-ISA placing UE authorized-only state
- UE sends DHCP Discover which is forwarded by the MS-ISA to the CPM
- CPM places UE in ESM-user state
- Upon disconnect, the AP sends a RADIUS accounting-stop which clears the UE context on the WLAN-GW

Figure 229 **Call Flow for Secure SSID with DSM**

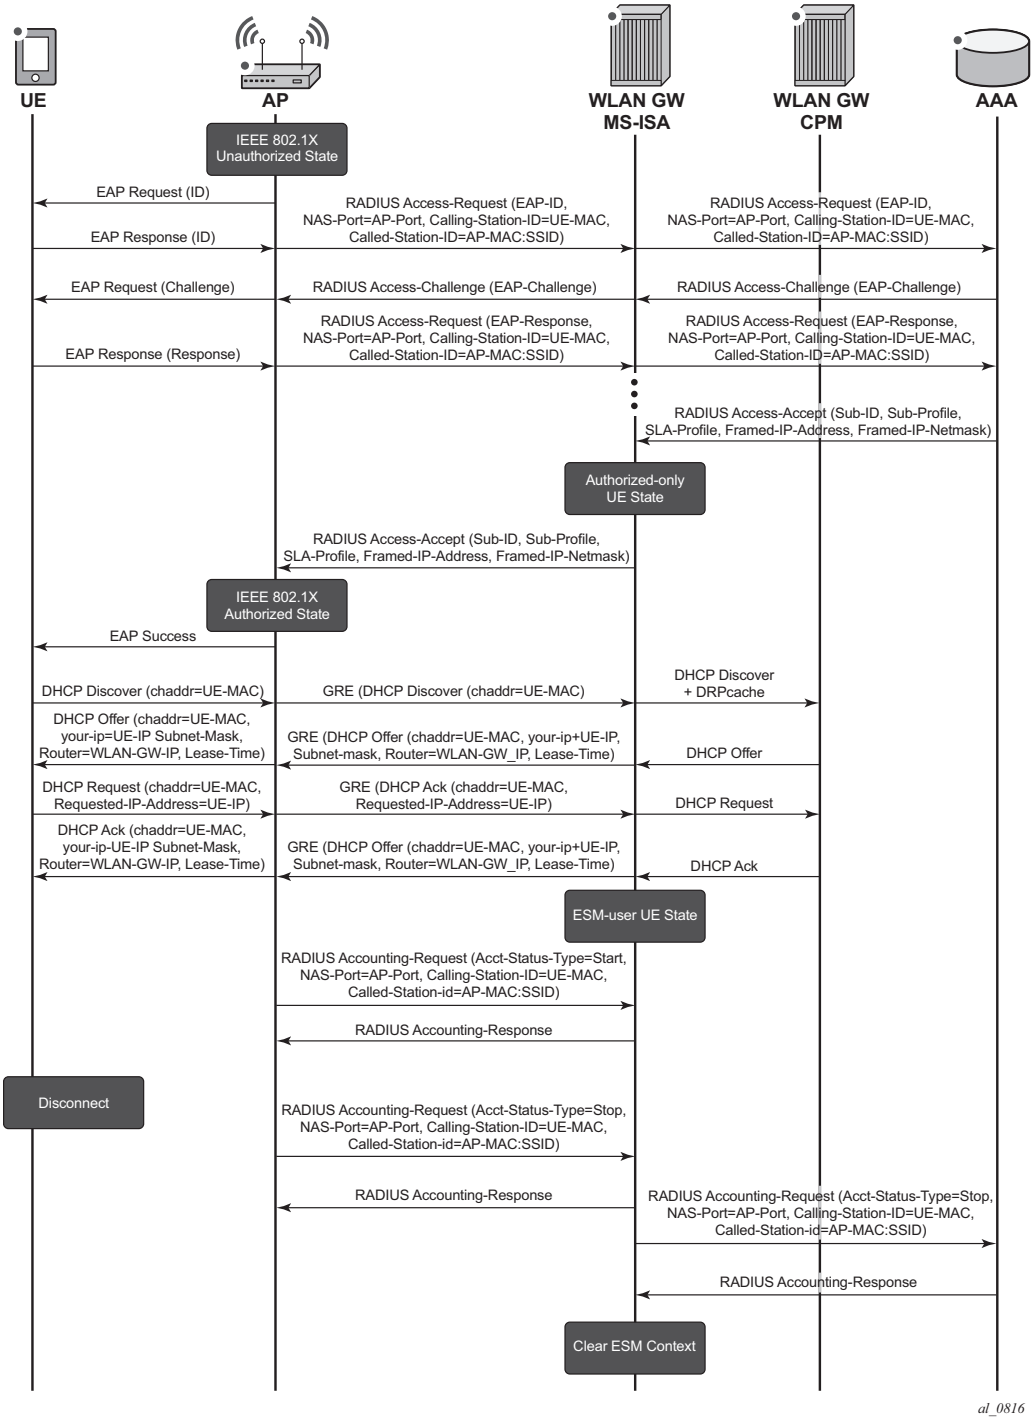

Debug

In this example, the following debug configuration is used (note that some default options are automatically added and do not need to be entered manually, e.g. mode under dhcp). For DRP, only a limited number of UEs can be debugged at a time and their MAC address have to be specified.

```
debug
    router "2001"
        ip
            dhcp
                detail-level medium
                mode egr-ingr-and-dropped
            exit
        exit
    exit
    wlan-gw
        group 1
            ue 68:7f:74:8b:3d:d7 packet radius dhcp
        exit
    exit
exit
```

The following is a partial debug of a UE connecting to the SSID and authenticating with the RADIUS server. Shortly after logging in the UE disconnects from the SSID and the subscriber is removed on reception of the RADIUS accounting-stop message.

As soon as the UE attempts to connect to the secure SSID, the WLAN-GW distributed RADIUS proxy in VPRN 1000 receives the first Access-Request packet from the AP. Note that the CALLING STATION ID [31] attribute contains the MAC address of the UE, and that the AP sends the SSID name in the NAS IDENTIFIER [32] attribute.

```
1464 2015/09/15 16:19:52.93 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 3291
   Info:     anchor ingressing frame
             radius upstream from client

   IP/UDP:   from 10.1.3.10 (port 51235) to 192.11.11.11 (port 1812)

   RADIUS:   Access-Request (1)  id 122  len 190
    USER NAME [1] 5 user1
    NAS IP ADDRESS [4] 4 10.1.3.10
    FRAMED IP ADDRESS [8] 4 255.255.255.255
    NAS IDENTIFIER [32] 10 WiFi-165-1
    CALLED STATION ID [30] 28 00-0D-67-39-0B-65:WiFi-165-1
    NAS PORT TYPE [61] 4 Wireless - IEEE 802.11(19)
    NAS PORT [5] 4 0
    CALLING STATION ID [31] 17 68-7F-74-8B-3D-D7
    CONNECT INFO [77] 21 CONNECT 0Mbps 802.11b
    SESSION ID [44] 17 556F2789-0000008D
    FRAMED MTU [12] 4 1400
```

```
    EAP MESSAGE [79] 10 0x02e2000a017573657231
    MESSAGE AUTHENTICATOR [80] 16 0xbc3a66d7f9d4e02465797f2018914ed7
"
```

The WLAN-GW MS-ISA forwards the Access-Request to the RADIUS server in the Base router.

```
1465 2015/09/15 16:19:52.94 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 3292
   Info:     anchor egressing frame
             radius upstream to server

   IP/UDP:   from 10.10.165.1 (port 1024) to 10.93.186.2 (port 1812)

   RADIUS:   Access-Request (1)  id 20  len 190
    USER NAME [1] 5 user1
    NAS IP ADDRESS [4] 4 10.1.3.10
    FRAMED IP ADDRESS [8] 4 255.255.255.255
    NAS IDENTIFIER [32] 10 WiFi-165-1
    CALLED STATION ID [30] 28 00-0D-67-39-0B-65:WiFi-165-1
    NAS PORT TYPE [61] 4 Wireless - IEEE 802.11(19)
    NAS PORT [5] 4 0
    CALLING STATION ID [31] 17 68-7F-74-8B-3D-D7
    CONNECT INFO [77] 21 CONNECT 0Mbps 802.11b
    SESSION ID [44] 17 556F2789-0000008D
    FRAMED MTU [12] 4 1400
    EAP MESSAGE [79] 10 0x02e2000a017573657231
    MESSAGE AUTHENTICATOR [80] 16 0xbf48919833584995109b8387efc03b21
"
```

Many RADIUS Access-Request and Access Challenge messages are exchanged, which encapsulate the EAP authentication between the UE and the RADIUS server. At the end of the exchange, for a successful authentication, the WLAN-GW receives an Access-Accept message (for a failed authentication it would receive an Access-Reject).

```
1506 2015/09/15 16:20:01.69 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 3333
   Info:     anchor ingressing frame
             radius downstream from server

   IP/UDP:   from 10.93.186.2 (port 1812) to 10.10.165.1 (port 1024)

   RADIUS:   Access-Accept (2)  id 25  len 203
    VSA [26] 6 Alcatel(6527)
      DEFAULT ROUTER [18] 4 192.168.1.1
    VSA [26] 6 Alcatel(6527)
      PRIMARY DNS [9] 4 10.92.186.92
    FRAMED IP ADDRESS [8] 4 192.168.1.9
    FRAMED IP NETMASK [9] 4 255.255.255.0
    USER NAME [1] 5 user1
    VSA [26] 52 Microsoft(311)
      MS MPPE RECV KEY [17] 50 0xc1af6befb148f03d5bd9bb8863500dd0a1ffcf57392dcda
8db5529be6e2de52fc239d3595212ee1b181e50c064e292595db8
    VSA [26] 52 Microsoft(311)
      MS MPPE SEND KEY [16] 50 0xcbe708a0751bc3c9ef43bb58e2b103cca0a6373b6800279
```

```
148a0f1934176f000e1540e5078eeba9d43af5f42d4799b16a79d
    EAP MESSAGE [79] 4 0x03ec0004
    MESSAGE AUTHENTICATOR [80] 16 0xb3d2459f830217fd455b26e7767012c3
"
```

The Access-Accept contains the IP addressing parameters for the UE such as the IP address, netmask, and default gateway, as well as the subscriber ID string. The IP addressing information is used by the WLAN-GW, but the Access-Accept message is also forwarded by the RADIUS proxy to the AP to tell it that the UE authenticated successfully so it can associate with the SSID.

```
1507 2015/09/15 16:20:01.69 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 3334
   Info:     anchor egressing frame
             radius downstream to client

   IP/UDP:   from 192.11.11.11 (port 1812) to 10.1.3.10 (port 51235)

   RADIUS:   Access-Accept (2)  id 132  len 203
    VSA [26] 6 Alcatel(6527)
      DEFAULT ROUTER [18] 4 192.168.1.1
    VSA [26] 6 Alcatel(6527)
      PRIMARY DNS [9] 4 10.92.186.92
    FRAMED IP ADDRESS [8] 4 192.168.1.9
    FRAMED IP NETMASK [9] 4 255.255.255.0
    USER NAME [1] 5 user1
    VSA [26] 52 Microsoft(311)
      MS MPPE RECV KEY [17] 50 0xc1afa8a2e9f23dbe5c0d41410a8bcc7fc42406813a3bff6
a61c957fbad58b7af6de0447898603980aeebe5cc2d5db54b8ca7
    VSA [26] 52 Microsoft(311)
      MS MPPE SEND KEY [16] 50 0xcbe7fb9182312534ea50ecdfc8ed59874401515968ae276
7826fa664e3871d0b13e2946b01750825dbb95b3fe6ee615afa1a
    EAP MESSAGE [79] 4 0x03ec0004
    MESSAGE AUTHENTICATOR [80] 16 0x66b269e340328cee108dfc1d27f46fca
"
```

After the AP allows the UE to connect to the secure SSID, establishing L2 connectivity to the WLAN-GW across the soft-GRE tunnel, the UE can obtain an IP address through DHCP. The WLAN-GW receives a DHCP Discover from the UE on MS-ISA MDA 2/1:

```
1508 2015/09/15 16:20:01.83 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 3335
   Info:     anchor ingressing frame
             received upstream from tunnel

   Ethernet: from 68:7f:74:8b:3d:d7 to ff:ff:ff:ff:ff:ff (ethertype: 0x0800)

   IP/UDP:   from 0.0.0.0 (port 68) to 255.255.255.255 (port 67)

   DHCP:
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 68:7f:74:8b:3d:d7    xid: 0xfb4fb37
```

```
   DHCP options:
   [53] Message type: Discover
   [61] Client id: (hex) 01 68 7f 74 8b 3d d7
   [12] Host name: VMS11
   [60] Class id: MSFT 5.0
   [55] Param request list: len = 12
             1  Subnet mask
            15  Domain name
             3  Router
             6  Domain name server
            44  NETBIOS name server
            46  NETBIOS type
            47  NETBIOS scope
            31  Router discovery
            33  Static route
           121  Unknown option
           249  Unknown option
            43  Vendor specific
   [255] End
"
```

The MS-ISA forwards the DHCP Discover to the CPM and it arrives on group interface GI1 in VPRN 2001.

```
1509 2015/09/15 16:20:01.83 EDT MINOR: DEBUG #2001 vprn2001 PIP
"PIP: DHCP
instance 6 (2001), interface index 3 (GI1),
   received DHCP Boot Request on Interface GI1 (2/1/nat-out-ip:2049.1) Port 67

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 68:7f:74:8b:3d:d7    xid: 0xfb4fb37

   DHCP options:
   [53] Message type: Discover
   [61] Client id: (hex) 01 68 7f 74 8b 3d d7
   [12] Host name: VMS11
   [60] Class id: MSFT 5.0
   [55] Param request list: len = 12
             1  Subnet mask
            15  Domain name
             3  Router
             6  Domain name server
            44  NETBIOS name server
            46  NETBIOS type
            47  NETBIOS scope
            31  Router discovery
            33  Static route
           121  Unknown option
           249  Unknown option
            43  Vendor specific
   [255] End
"
```

The WLAN-GW sends a DHCP Offer to the UE with the IP address information retrieved from the RADIUS Access-Accept message.

```
1510 2015/09/15 16:20:01.85 EDT MINOR: DEBUG #2001 vprn2001 PIP
"PIP: DHCP
instance 6 (2001), interface index 3 (GI1),
   transmitted DHCP Boot Reply to Interface GI1 (2/1/nat-out-ip:2049.1) Port 68

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 192.168.1.9
   siaddr: 192.168.1.1       giaddr: 192.168.1.1
   chaddr: 68:7f:74:8b:3d:d7    xid: 0xfb4fb37

   DHCP options:
   [53] Message type: Offer
   [54] DHCP server addr: 192.168.1.1
   [51] Lease time: 604800
   [1] Subnet mask: 255.255.255.0
   [3] Router: 192.168.1.1
   [6] Domain name server: 10.92.186.92
   [28] Broadcast addr: 192.168.1.255
   [61] Client id: (hex) 01 68 7f 74 8b 3d d7
   [12] Host name: VMS11
   [255] End
"
```

The Offer message is sent to the MS-ISA and towards the UE (not shown). The UE then sends a DHCP Request and the WLAN-GW responds with an Ack.

```
1513 2015/09/15 16:20:01.86 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 3337
   Info:     anchor ingressing frame
             received upstream from tunnel

   Ethernet: from 68:7f:74:8b:3d:d7 to ff:ff:ff:ff:ff:ff (ethertype: 0x0800)

   IP/UDP:   from 0.0.0.0 (port 68) to 255.255.255.255 (port 67)

   DHCP:
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 68:7f:74:8b:3d:d7    xid: 0xfb4fb37

   DHCP options:
   [53] Message type: Request
   [61] Client id: (hex) 01 68 7f 74 8b 3d d7
   [50] Requested IP addr: 192.168.1.9
   [54] DHCP server addr: 192.168.1.1
   [12] Host name: VMS11
   [81] client FQDN: rcode1: 0, rcode2: 0, domain name = (hex) 00 56 4d 53 31
   31
   [60] Class id: MSFT 5.0
   [55] Param request list: len = 12
             1  Subnet mask
            15  Domain name
             3  Router
             6  Domain name server
            44  NETBIOS name server
            46  NETBIOS type
            47  NETBIOS scope
            31  Router discovery
```

```
            33  Static route
           121  Unknown option
           249  Unknown option
            43  Vendor specific
   [255] End
"

1515 2015/09/15 16:20:01.86 EDT MINOR: DEBUG #2001 vprn2001 PIP
"PIP: DHCP
instance 6 (2001), interface index 3 (GI1),
   transmitted DHCP Boot Reply to Interface GI1 (2/1/nat-out-ip:2049.1) Port 68

   H/W Type: Ethernet(10Mb)  H/W Address Length: 6
   ciaddr: 0.0.0.0           yiaddr: 192.168.1.9
   siaddr: 192.168.1.1       giaddr: 192.168.1.1
   chaddr: 68:7f:74:8b:3d:d7    xid: 0xfb4fb37

   DHCP options:
   [53] Message type: Ack
   [54] DHCP server addr: 192.168.1.1
   [51] Lease time: 604800
   [1] Subnet mask: 255.255.255.0
   [3] Router: 192.168.1.1
   [6] Domain name server: 10.92.186.92
   [28] Broadcast addr: 192.168.1.255
   [61] Client id: (hex) 01 68 7f 74 8b 3d d7
   [12] Host name: VMS11
   [81] client FQDN: rcode1: 0, rcode2: 0, domain name = (hex) 00 56 4d 53 31
   31
   [255] End
"
```

The AP sends a RADIUS accounting Start to the WLAN-GW as a result of the UE successfully associating with the SSID.

```
1518 2015/09/15 16:20:01.88 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 3339
   Info:     anchor ingressing frame
             radius upstream from client

   IP/UDP:   from 10.1.3.10 (port 51236) to 192.11.11.11 (port 1813)

   RADIUS:   Accounting-Request (4)  id 133  len 197
    SESSION ID [44] 17 556F2789-0000008D
    EVENT TIMESTAMP [55] 4 1442292269
    STATUS TYPE [40] 4 Start(1)
    AUTHENTIC [45] 4 RADIUS(1)
    USER NAME [1] 5 user1
    NAS IP ADDRESS [4] 4 10.1.3.10
    FRAMED IP ADDRESS [8] 4 192.168.1.9
    NAS IDENTIFIER [32] 10 WiFi-165-1
    CALLED STATION ID [30] 28 00-0D-67-39-0B-65:WiFi-165-1
    NAS PORT TYPE [61] 4 Wireless - IEEE 802.11(19)
    NAS PORT [5] 4 0
    CALLING STATION ID [31] 17 68-7F-74-8B-3D-D7
    CONNECT INFO [77] 21 CONNECT 0Mbps 802.11b
    SESSION ID [44] 17 556F2789-0000008D
    DELAY TIME [41] 4 0
```

```
"
```

At the end of the session, the UE disconnects from the SSID, and the AP sends a RADIUS accounting Stop to the WLAN-GW.

```
1520 2015/09/15 16:20:37.45 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 3401
   Info:     anchor ingressing frame
             radius upstream from client

   IP/UDP:   from 10.1.3.10 (port 51237) to 192.11.11.11 (port 1813)

   RADIUS:   Accounting-Request (4)  id 134  len 233
    SESSION ID [44] 17 556F2789-0000008D
    EVENT TIMESTAMP [55] 4 1442292304
    STATUS TYPE [40] 4 Stop(2)
    AUTHENTIC [45] 4 RADIUS(1)
    USER NAME [1] 5 user1
    NAS IP ADDRESS [4] 4 10.1.3.10
    FRAMED IP ADDRESS [8] 4 192.168.1.9
    NAS IDENTIFIER [32] 10 WiFi-165-1
    CALLED STATION ID [30] 28 00-0D-67-39-0B-65:WiFi-165-1
    NAS PORT TYPE [61] 4 Wireless - IEEE 802.11(19)
    NAS PORT [5] 4 0
    CALLING STATION ID [31] 17 68-7F-74-8B-3D-D7
    CONNECT INFO [77] 21 CONNECT 0Mbps 802.11b
    SESSION ID [44] 17 556F2789-0000008D
    DELAY TIME [41] 4 0
    SESSION TIME [46] 4 35
    INPUT PACKETS [47] 4 57
    OUTPUT PACKETS [48] 4 36
    INPUT OCTETS [42] 4 5228
    OUTPUT OCTETS [43] 4 5538
    TERMINATE CAUSE [49] 4 User Request(1)
"
```

This removes the subscriber from the WLAN-GW, clears the DHCP state, and also removes the GRE tunnel if this UE is the last one on the tunnel.

Conclusion

The 7750 SR WLAN-GW can support many WiFi Offload architectures including secure SSID with various types of EAP authentication. WiFi Offload functions such as terminating GRE tunnels, NAT, and RADIUS server connectivity can be performed in separate routing instances if required. UE IP addresses can be assigned locally or from an external source such as RADIUS, and routing to the Internet can be performed in various ways, including NAT. Several show commands and debug options are available to help the operator monitor and troubleshoot the solution.

WiFi Aggregation and Offload — IPv4/v6 Dual-Stack UEs

This chapter provides information about WiFi aggregation and offload IPv4/v6 dual-stack UEs.

Topics in this chapter include:

Applicability

This chapter is applicable to all 7750 SR platforms that function as WLAN gateways (WLAN-GWs). The configuration makes use of soft Generic Routing Encapsulation (GRE) as the access technology, which requires one or more WLAN-IOMs (IOM3-XP and 2 x MS-ISAs). However, a similar configuration can be made without soft-GRE and, therefore, without WLAN-IOMs.

The example configuration was tested on a 7750 SR-7 with SR OS 13.0.R3.

Summary

WiFi Aggregation and Offload functionality for the 7750 SR is supported on SR OS 10.0.R3 or later. The functionality includes enhanced subscriber management (ESM) for user equipment (UE) that gains network access via a WiFi service.

This chapter provides a functional description of the WLAN-GW features related to IPv4/v6 dual-stack UEs, as well as the related configuration.

Overview

Because IP address demand is mainly due to mobile devices, the support of IPv6 on mobile devices is a major requirement to manage IPv4 address depletion.

However, IPv6 on mobile devices is currently considered as an add-on rather than a replacement of IPv4, so the demand is for IPv4/v6 dual-stack UEs.

The basic concepts of ESMv6 for IPoE dual-stack hosts also apply to dual-stack UEs. However, WLAN-GW operates in a bridged environment where the Access Point (AP) performs L2 forwarding of Ethernet frames between the IEEE 802.11 air interface and the soft-GRE, soft-L2TPv3, or VLAN tunnel. Therefore, a WLAN-GW treats each UE as an individual subscriber who connects to the WiFi service. This contrasts with ESMv6 IPoE hosts behind a routed residential gateway (RG), where multiple hosts connect via the RG and the BNG treats the RG as the subscriber.

Depending on the type of UE, it may be allocated an IPv4 address through DHCPv4 and an IPv6 address through Stateless Address Auto-Configuration (SLAAC), or DHCPv6, or both (not all UEs have support for DHCPv6). Therefore, a UE can instantiate up to three IPoE hosts: a DHCPv4 host, a SLAAC host, and a DHCPv6 host.

Authentication and authorization depend on whether the UE connects to a WiFi with open or closed SSID. With an open SSID, authentication and authorization are performed when the first packet is received from the UE (typically a DHCPv4 Discover, an ICMPv6 Router-Solicit, or a DHCPv6 Solicit), similar to the routed RG model. Upon successful authentication, the Access-Accept is stored for 10 s on the WLAN-GW, so for a dual-stack IPv4/v6 UE, two or three authentication rounds can be avoided if DHCPv4, SLAAC, and DHCPv6 are started within this 10 s interval.

When the UE has successfully authenticated with the portal, a CoA-Request may lift the HTTP redirect filter by changing the SLA profile and, optionally, the subscriber profile. If the CoA-Request contains the subscriber ID, the CoA-Request applies to both the DHCPv4 host and the SLAAC and/or DHCPv6 host. See the RADIUS attributes reference guide for more information about alternative subscriber host identification in RADIUS CoA-Request messages.

Figure 230 **DHCPv4 + SLAAC/64 — Open SSID**

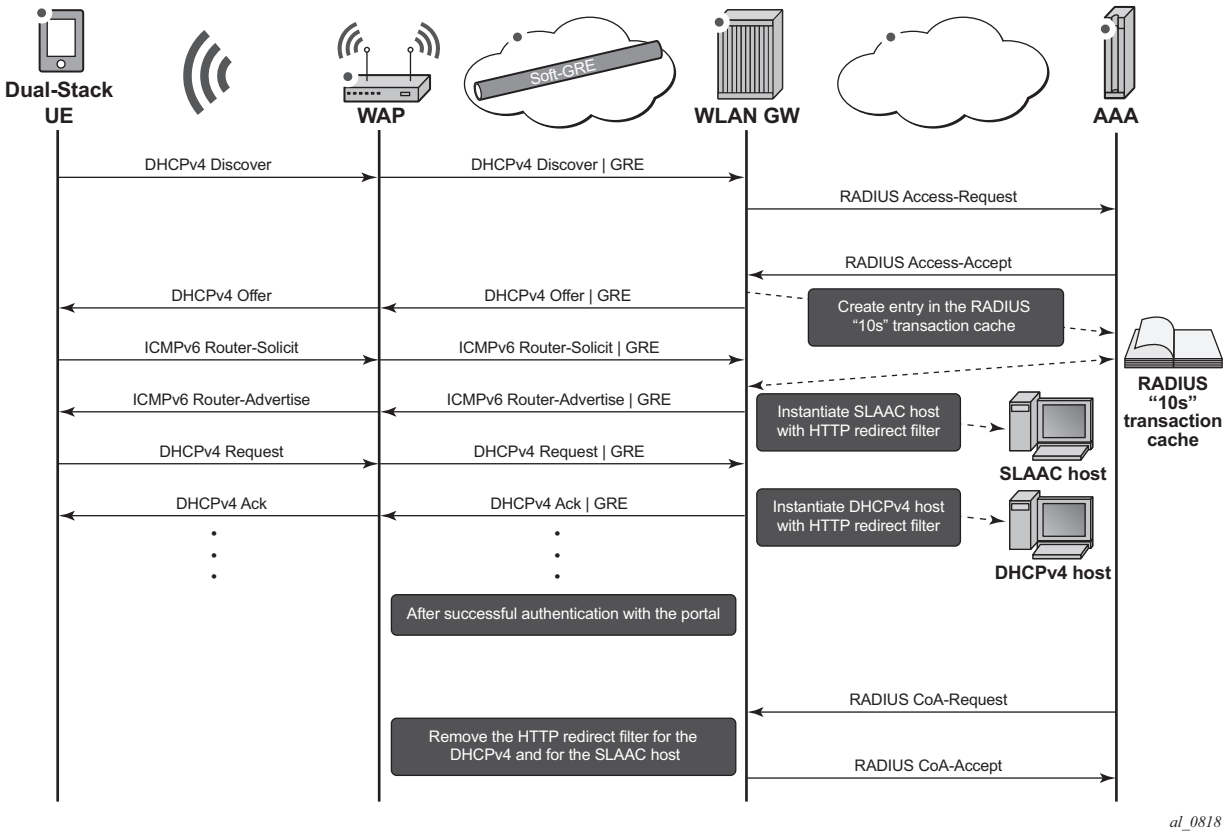

With a closed SSID, there is a separation between the authentication and authorization phases. When a UE connects to a WiFi with closed SSID in WPA-Enterprise mode, also known as WPA-802.1X mode, the UE initiates authentication before it obtains an IP address. The WLAN-GW is aware of the successful authentication when it receives the DHCPv4 Discover, the ICMPv6 Router-Solicit, or the DHCPv6 Solicit.

As with ESMv6, the WLAN-GW supports SLAAC/64 and DHCPv6/128 Identity Association for Non-temporary Addresses (IA_NA). DHCPv6 Identity Association for Prefix Delegation (IA_PD) is not supported because the UEs are considered as individual hosts that have direct Layer 2 connectivity with the WLAN-GW. Devices that use the UE as an IPv6 gateway are currently not supported.

For SLAAC/64 hosts, the DNS information can be advertised with the recursive DNS server (RDNSS) option [RFC6106] via SLAAC or via stateless DHCPv6 [RFC3736]. For DHCPv6/128 hosts, the DNS information is advertised via DNS options for DHCPv6 [RFC3646]. If the AP supports a Lightweight DHCPv6 Relay Agent (LDRA), the WLAN-GW can learn the AP MAC address and the SSID that the UE connects to if the DHCPv6 Interface-Id option is in the format **<ap-mac>:<ssid>:{o (open) | s (secure)}** This information can then be used in subsequent accounting messages.

The following three IPv4/v6 dual-stack UE IP address assignment models are available:

- DHCPv4 + SLAAC/64
- DHCPv4 + SLAAC/64 with DHCPv4 linking
- DHCPv4 + DHCPv6/128 IA_NA

In the DHCPv4 + SLAAC/64 model, DHCPv4 DORA and SLAAC/64 are processed independently of each other. If successful, two IPoE hosts are instantiated on the WLAN-GW for a particular UE: a DHCPv4 IPoE host and an IPv6 SLAAC/64 host.

Figure 231 **DHCPv4 + SLAAC/64 Model — Closed SSID**

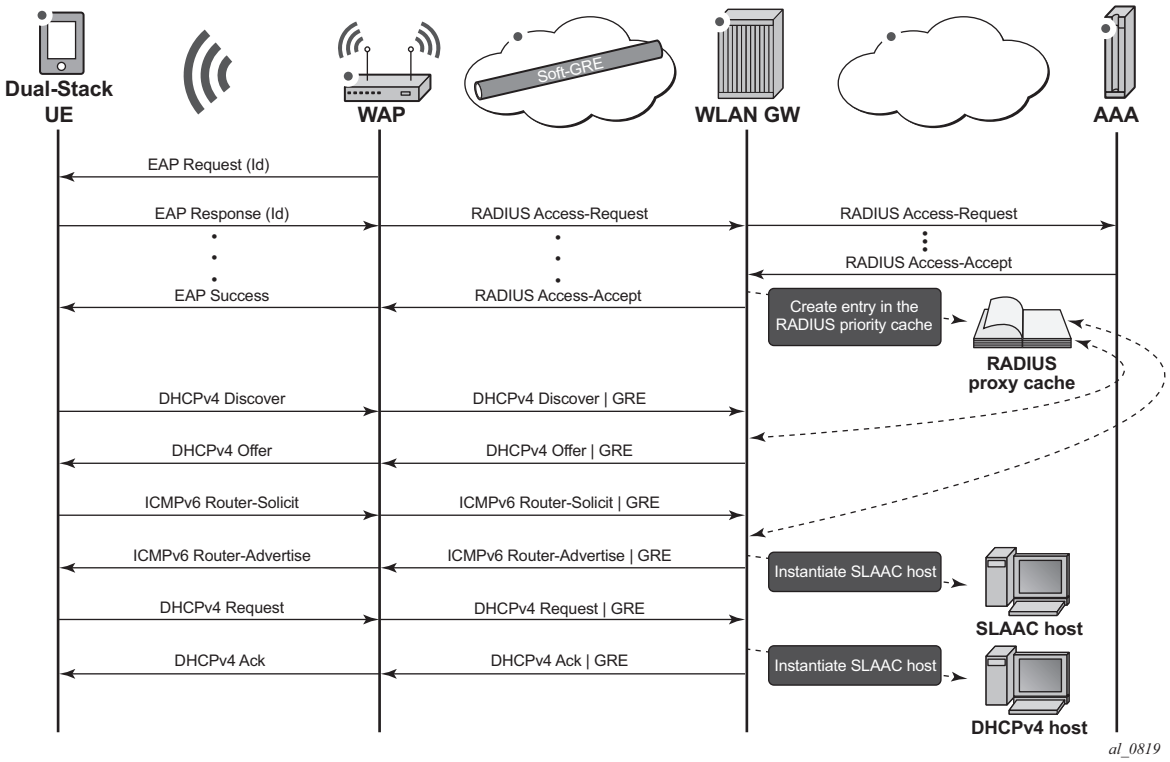

When the AP sends a RADIUS Accounting-Stop for a particular UE while track-accounting is enabled for Accounting-Stop messages, both the DHCPv4 IPoE host and the IPv6 SLAAC/64 host will be removed.

However, it is not always possible for the AP to send RADIUS accounting messages (for example, in the case of an open SSID). Because SLAAC has no renew or release mechanism, the only way to delete a SLAAC host is to determine which UE was stopped using the SLAAC prefix; for example, by using idle-timeout and/or by periodic Subscriber Host Connectivity Verification (SHCV).

In the DHCPv4 + SLAAC/64 with DHCPv4 linking model, a SLAAC/64 host is instantiated when a DHCPv4 host is instantiated. The state of the SLAAC/64 host is linked to the state of the DHCPv4 host. This is useful to speed up the removal of the SLAAC host in cases where the AP does not send RADIUS accounting messages. With DHCPv4 linking, when the DHCPv4 host is removed, also the SLAAC/64 host is removed.

Figure 232 **DHCPv4 + SLAAC/64 with DHCPv4 Linking Model — Closed SSID**

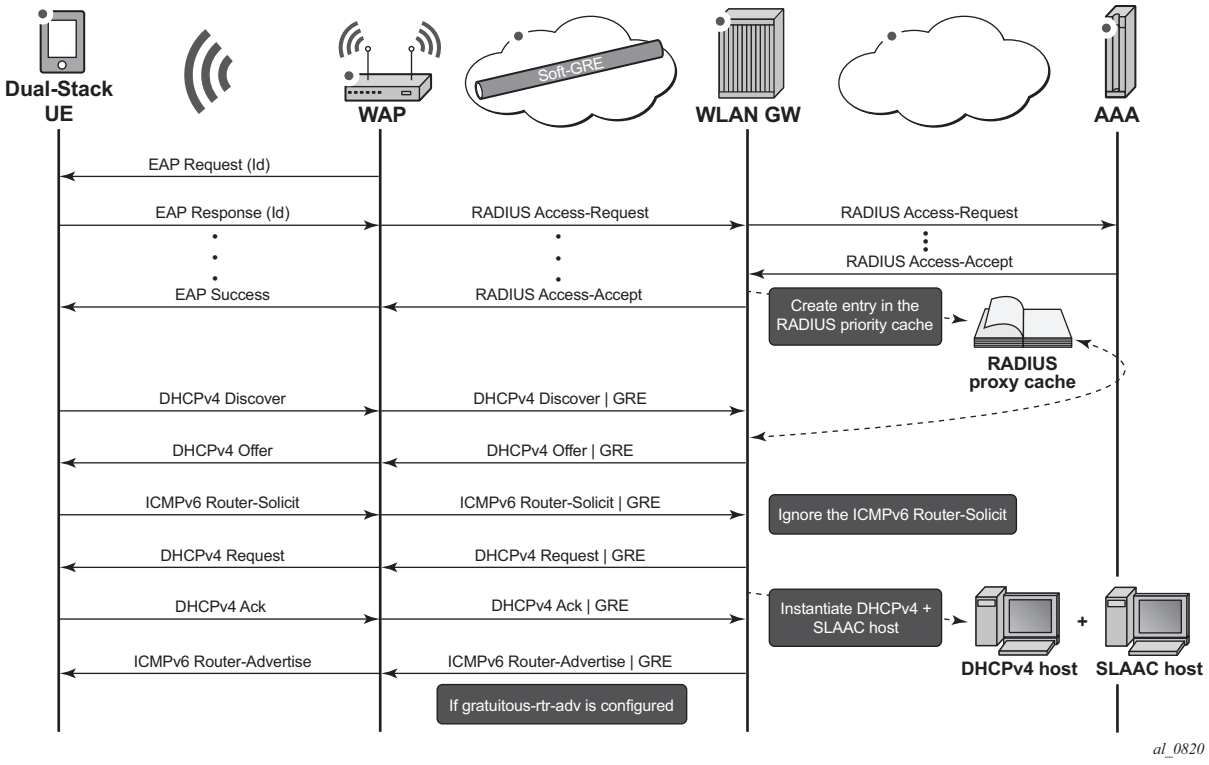

In the DHCPv4 + DHCPv6/128 IA_NA model, similar to the DHCPv4 + SLAAC/64 model, DHCPv4 DORA and DHCPv6/128 IA_NA are processed independently of each other. SLAAC/64 is optional in this model although it is typically enabled because some UEs do not support DHCPv6.

UEs that do support stateful address auto-configuration only initiate DHCPv6 when they receive an ICMPv6 Router-Advertisement with the M-bit set (RFC 2462, *IPv6 Stateless Address Autoconfiguration*). Because the WLAN-GW does not know whether the UE supports DHCPv6, the WLAN-GW must include a SLAAC/64 prefix in the ICMPv6 Router-Advertisement, also for UEs that do support DHCPv6. Therefore, for a UE that does support DHCPv6, three IPoE hosts are instantiated in the WLAN-GW. To avoid this, the WLAN-GW can be configured to flush the SLAAC/64 host when a DHCPv6/128 IA_NA host is established. The UE should always prefer the DHCPv6/128 IA_NA address for sending data traffic above the IPv6 address derived from the SLAAC/64 prefix.

Figure 233 **DHCPv4 + DHCPv6/128 IA_NA Model — Closed SSID**

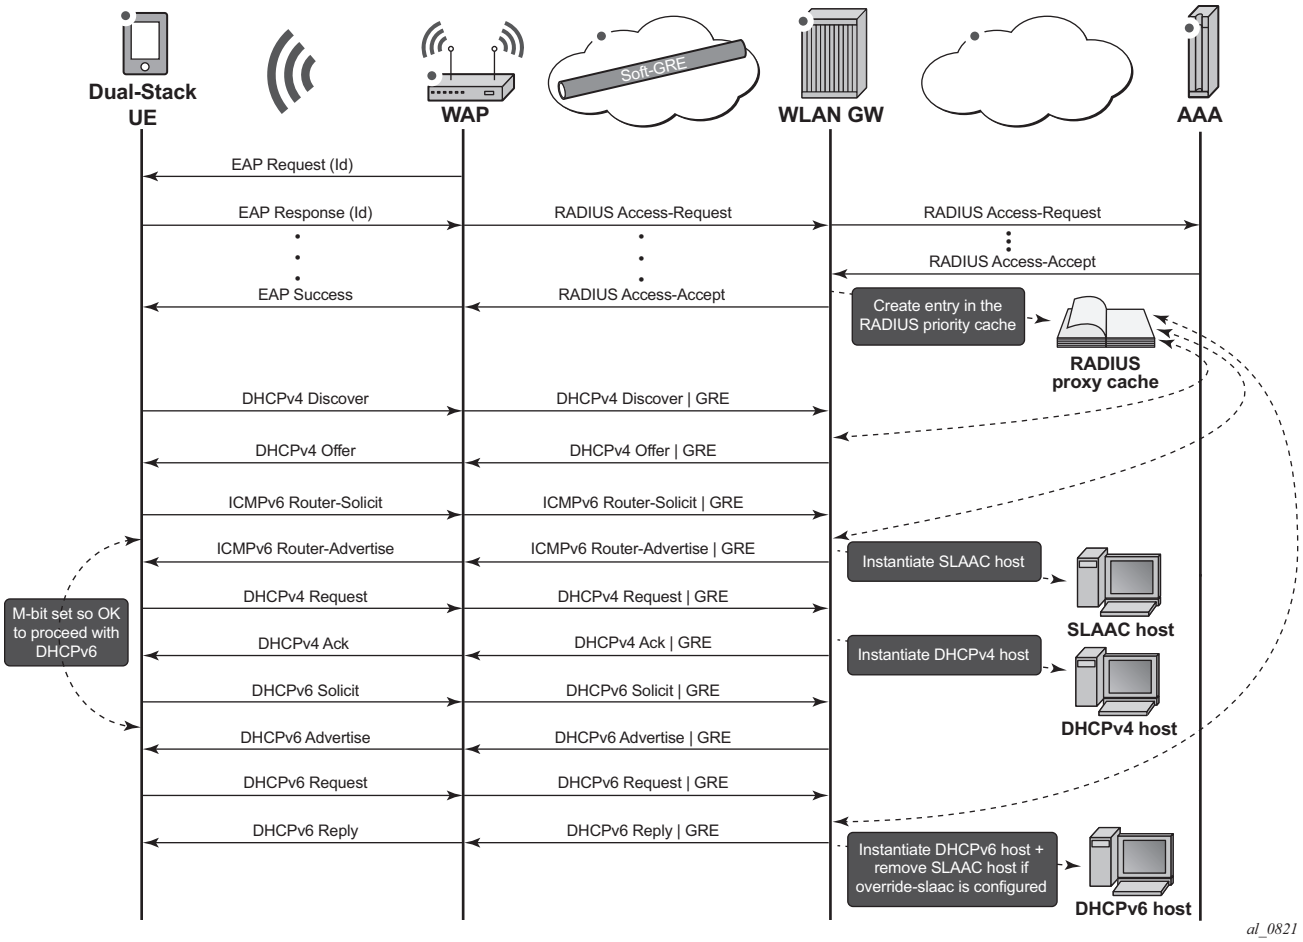

As with ESMv6, the SLAAC/64 prefix could come from the Local User Database (LUDB); this is typically not used because it requires configuring individual UE MAC addresses with their associated SLAAC/64 prefix. Alternatively, the SLAAC/64 prefix could come from RADIUS via the Framed-IPv6-Prefix attribute, or from a local SLAAC prefix pool that is referenced in the LUDB or from RADIUS via the Alc-SLAAC-IPv6-Pool attribute.

The DHCPv6/128 prefix comes from a DHCPv6 server that could be local (collocated with the WLAN-GW) or external, or from RADIUS via the Alc-IPv6-Address attribute. When a DHCPv6 server is used, the WLAN-GW relays the DHCPv6 messages between the UE and the local or external DHCPv6 server. If the DHCPv6/128 prefix comes from RADIUS/LUDB, the WLAN-GW must be configured as a DHCPv6 proxy server.

Note that IPv6 for WLAN-GW UEs is not supported in combination with certain other features, which include GPRS Tunneling Protocol (GTP(v2)) offload, migrant UEs, and data-triggered authentication (DTA).

Data-triggered authentication is not supported for IPv6 hosts, which means that an IPv6 packet from a UE for which no ESM context exists will not trigger RADIUS authentication. However, by using SLAAC/64 with DHCPv4 linking, the SLAAC host will be created together with the DHCPv4 host by successful completion of IPv4 data-triggered authentication. This requires the RADIUS Access-Accept to contain the necessary DHCPv4 and SLAAC/64 attributes.

Figure 234 **DHCPv4 + SLAAC/64 with DHCPv4 Linking Model — DTA**

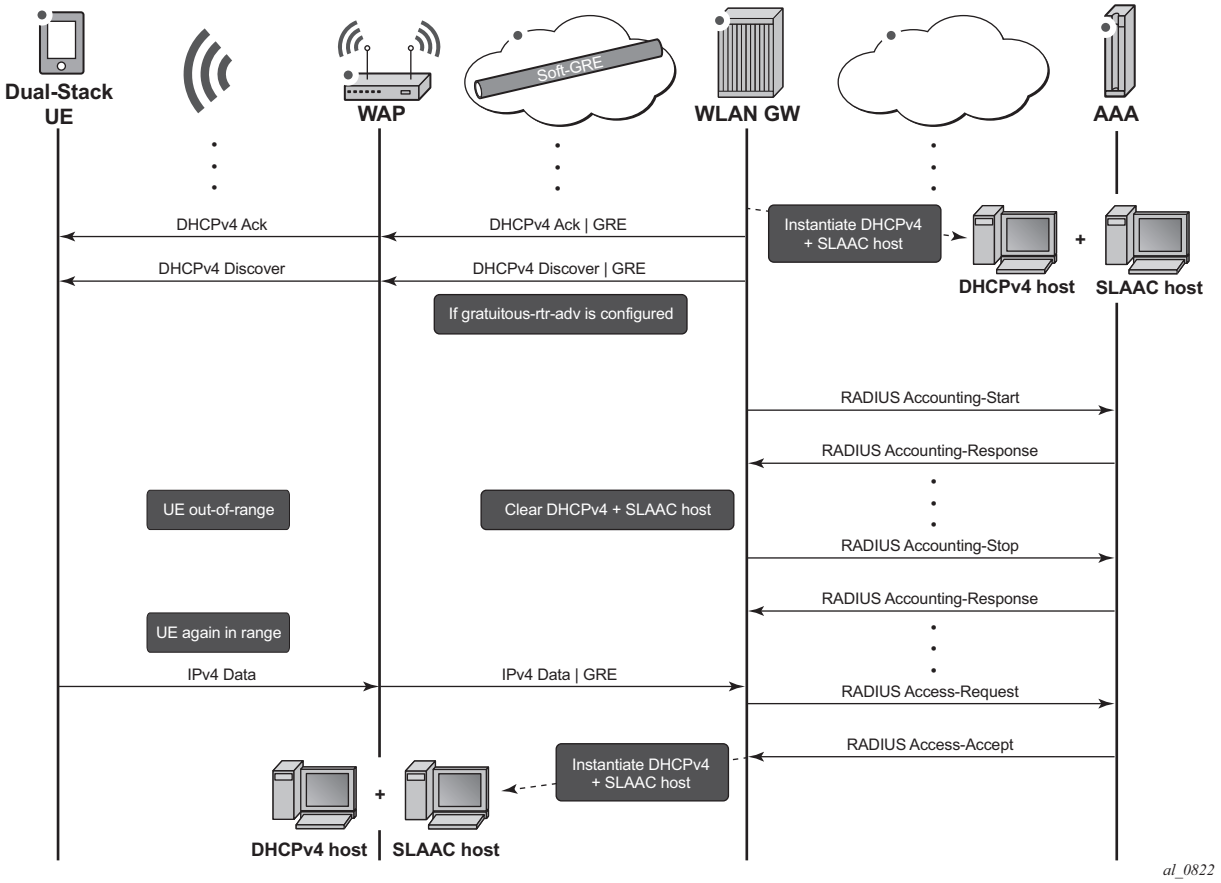

IPv6 is also not supported for migrant UEs, which means that ICMPv6 Router-Solicitation and DHCPv6 Solicit messages will be dropped by the WLAN-GW as long as the UE is in a migrant state. However, by using SLAAC/64 with DHCPv4 linking, when the UE becomes an ESM subscriber and a DHCPv4 host is created, a SLAAC/64 host is also created.

Configuration

Open Versus Closed SSID

The configuration examples in this chapter always refer to a closed SSID scenario. With an open SSID, the lookup in the RADIUS proxy cache is typically not configured. Instead, an authentication policy is directly referenced.

```
configure service vprn 2 customer 1 create
    subscriber-interface "sub-int-1" create
        group-interface "group-int-1" wlangw create
            authentication-policy "auth-pol-1"
            dhcp
                no user-db
            exit
        exit
    exit
exit
```

DHCPv4 + SLAAC/64 Model

In this model, DHCPv4 and SLAAC/64 are enabled independently of each other. The autonomous flag tells the UE that the IPv6 prefix in the ICMPv6 Router-Advertisement can be used for SLAAC. The no on-link configuration commands the UE to always perform neighbor discovery for the WLAN-GW, even for destinations within the IPv6 prefix.

```
configure service vprn 2 customer 1 create
    subscriber-interface "sub-int-1" create
        address 10.255.255.254/8
        ipv6
            subscriber-prefixes
                prefix 2001:db8:ffff::/48 wan-host
            exit
        exit
        group-interface "group-int-1" wlangw create
            ipv6
                router-advertisements
                    no managed-configuration
                    no other-stateful-configuration
                    dns-options
                        include-dns
                    exit
                    prefix-options
                        autonomous
                        no on-link
                    exit
                    no shutdown
                exit
                router-solicit
                    user-db "ludb-1"
                    no shutdown
                exit
            exit
            ipoe-linking
                shutdown
            exit
            sap-parameters
                sub-sla-mgmt
                    def-sub-id use-auto-id
                    sub-ident-policy "policy-sub-ident-1"
                exit
```

```
            exit
            dhcp
                proxy-server
                    emulated-server 172.16.0.1
                    no shutdown
                exit
                lease-populate 10000
                user-db "ludb-1"
                no shutdown
            exit
            ip-mtu 1454
            wlan-gw
                gw-address 172.16.74.244
                gw-ipv6-address 2001:db8::1:1
                router 1
                tcp-mss-adjust 1400
                wlan-gw-group 1
                no shutdown
            exit
        exit
    exit
    no shutdown
exit
```

The SLAAC/64 prefix can come from the RADIUS server, as in the following RADIUS users file:

```
"user-1" Cleartext-Password := "pass-1"
        Alc-Subsc-ID-Str := "user-1",
        Alc-Subsc-Prof-Str := "sub-profile-1",
        Alc-SLA-Prof-Str := "sla-profile-1",
        Framed-IP-Address := 10.255.0.1,
        Alc-Primary-DNS := 67.138.54.100,
        Framed-IPv6-Prefix := 2001:db8:ffff::/64,
        Alc-IPv6-Primary-Dns := 2001:db8::8:8:8:8,
        Alc-IPv6-Secondary-Dns := 2001:db8::8:8:4:4
```

If the UE is successfully connected, two IPoE hosts will exist on the WLAN-GW.

```
*A:WLAN-GW # show service active-subscribers
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber user-1 (sub-profile-1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[4/2/nat-out-ip:2049.4] - sla:sla-profile-1
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.255.0.1
                b0:9f:ba:b9:40:f8 N/A       DHCP
2001:db8:ffff::/64
                b0:9f:ba:b9:40:f8 N/A       IPoE-SLAAC
-------------------------------------------------------------------------------
```

```
Number of active subscribers : 1
-------------------------------------------------------------------------------
```

The trigger that created the SLAAC host and the origin is shown by issuing:

```
*A:WLAN-GW # show service id 2 slaac host detail
===============================================================================
SLAAC hosts for service 2
===============================================================================
Service ID           : 2
Prefix               : 2001:db8:ffff::/64
Interface Id         : N/A
Mac Address          : b0:9f:ba:b9:40:f8
Subscriber-interface : sub-int-1
Group-interface      : group-int-1
SAP                  : [4/2/nat-out-ip:2049.4]
Creation Time        : 2015/07/09 11:24:19
Persistence Key      : N/A

IPoE Trigger         : rtr-solicit
Last Auth Time       : 2015/07/09 11:24:19
Inactivity Timer     : 0d 00:03:59

Sub-Ident            : "user-1"
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-1"
App-Profile-String   : ""
ANCP-String          : ""
Int Dest Id          : ""
Category-Map-Name    : ""

Info origin          : radius
Pool                 : ""

Primary-Dns          : 2001:db8::8:8:8:8
Secondary-Dns        : 2001:db8::8:8:4:4

Circuit Id           : N/A
Remote Id            : N/A
-------------------------------------------------------------------------------
Number of hosts : 1
===============================================================================
```

The SLAAC/64 prefix can also come from a local SLAAC prefix pool:

```
configure service vprn 2 customer 1 create
    dhcp6
        local-dhcp-server "local-dhcp-server-1" create
            use-pool-from-client
            pool "slaac-prefix-pool-1" create
                prefix 2001:db8:ffff:ffff::/64 wan-host create
                    options
                        dns-server 2001:db8::8:8:8:8
                    exit
                exit
            exit
```

```
            no shutdown
        exit
    exit
exit
```

The subscriber interface must then be configured with local-address-assignment enabled:

```
configure service vprn 2 customer 1 create
    subscriber-interface "sub-int-1" create
        group-interface "group-int-1" wlangw create
            local-address-assignment
                ipv6
                    client-application ipoe-slaac
                    server "local-dhcp-server-1"
                exit
                no shutdown
            exit
        exit
    exit
exit
```

The origin of the SLAAC host then changes to:

```
*A:WLAN-GW # show service id 2 slaac host detail | match origin
Info origin          : localPool
```

DHCPv4 + SLAAC/64 with DHCPv4 linking model

In this model, DHCPv4 linking instantiates a SLAAC/64 host when a DHCPv4 host is instantiated. This requires **ipoe-linking** to be configured:

```
configure service vprn 2 customer 1 create
    subscriber-interface "sub-int-1" create
        address 10.255.255.254/8
        ipv6
            subscriber-prefixes
                prefix 2001:db8:ffff::/48 wan-host
            exit
        exit
        group-interface "group-int-1" wlangw create
            ipv6
                router-advertisements
                    no managed-configuration
                    no other-stateful-configuration
                    dns-options
                        include-dns
                    exit
                    prefix-options
                        autonomous
                        no on-link
                    exit
                    no shutdown
                exit
                router-solicit
                    shutdown
```

```
                exit
            exit
            ipoe-linking
                gratuitous-rtr-adv
                no shutdown
            exit
            sap-parameters
                sub-sla-mgmt
                    def-sub-id use-auto-id
                    sub-ident-policy "policy-sub-ident-1"
                exit
            exit
            dhcp
                proxy-server
                    emulated-server 172.16.0.1
                    no shutdown
                exit
                lease-populate 10000
                user-db "ludb-1"
                no shutdown
            exit
            ip-mtu 1454
            wlan-gw
                gw-address 172.16.74.244
                gw-ipv6-address 2001:db8::1:1
                router 1
                tcp-mss-adjust 1400
                wlan-gw-group 1
                no shutdown
            exit
        exit
    exit
    no shutdown
exit
```

Note that DHCPv4 linking is mutually exclusive with ICMPv6 Router-Solicit handling. Configuring DHCPv4 linking while ICMPv6 Router-Solicit handling is still enabled results in the following error:

```
*A:WLAN-GW # configure service vprn 2 subscriber-interface "sub-int-1" group-
interface "group-int-1" ipoe-linking no shutdown
MINOR: SVCMGR #1543 Can't enable linking if router solicit authentication is enabled
```

Similarly, enabling ICMPv6 Router-Solicit handling while DHCPv4 linking is still enabled, results in the following error:

```
*A:WLAN-GW # configure service vprn 2 subscriber-interface "sub-int-1" group-
interface "group-int-1" ipv6 router-solicit no shutdown
MINOR: SVCMGR #1544 Can't enable router solicit authentication if linking is enabled
```

As with the DHCPv4 + SLAAC/64 model without DHCPv4 linking, if the UE is successfully connected, two IPoE hosts will exist on the WLAN-GW:

```
*A:WLAN-GW # show service active-subscribers
```

```
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber user-1 (sub-profile-1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[4/2/nat-out-ip:2049.4] - sla:sla-profile-1
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.255.0.1
                b0:9f:ba:b9:40:f8 N/A       DHCP
2001:db8:ffff::/64
                b0:9f:ba:b9:40:f8 N/A       IPoE-SLAAC
-------------------------------------------------------------------------------
Number of active subscribers : 1
-------------------------------------------------------------------------------
```

The trigger that created the SLAAC host and the origin is shown by issuing:

```
*A:WLAN-GW # show service id 2 slaac host detail
===============================================================================
SLAAC hosts for service 2
===============================================================================
Service ID           : 2
Prefix               : 2001:db8:ffff::/64
Interface Id         : N/A
Mac Address          : b0:9f:ba:b9:40:f8
Subscriber-interface : sub-int-1
Group-interface      : group-int-1
SAP                  : [4/2/nat-out-ip:2049.4]
Creation Time        : 2015/07/09 11:49:42
Persistence Key      : N/A

IPoE Trigger         : linking
Last Auth Time       : N/A
Inactivity Timer     : N/A

Sub-Ident            : "user-1"
Sub-Profile-String   : "sub-profile-1"
SLA-Profile-String   : "sla-profile-1"
App-Profile-String   : ""
ANCP-String          : ""
Int Dest Id          : ""
Category-Map-Name    : ""

Info origin          : radius
Pool                 : ""

Primary-Dns          : 2001:db8::8:8:8:8
Secondary-Dns        : 2001:db8::8:8:4:4

Circuit Id           : N/A
Remote Id            : N/A
-------------------------------------------------------------------------------
Number of hosts : 1
```

```
===============================================================================
```

Clearing the DHCPv4 host results in both the DHCPv4 host and the SLAAC host being deleted.

```
*A:WLAN-GW # clear service id 2 dhcp lease-state mac b0:9f:ba:b9:40:f8

*A:WLAN-GW # show service active-subscribers
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
No active subscribers found
-------------------------------------------------------------------------------
```

DHCPv4 + DHCPv6/128 IA_NA Model

Because some UEs do not support DHCPv6, this model configures DHCPv4 + DHCPv6/128 IA_NA with SLAAC/64 enabled. To avoid having two IPv6oE hosts set up for the UEs that do support DHCPv6, the allo**w-multiple-wan-addr**esses and **override-slaac** parameters are both configured. The allow-multiple-wan-addresses allows handling of DHCPv6 when a SLAAC host exists already, and the **override-slaac** parameter removes the SLAAC host after successful assignment of an IPv6 address via DHCPv6:

```
configure service vprn 2 customer 1 create
    subscriber-interface "sub-int-1" create
        address 10.255.255.254/8
        ipv6
            subscriber-prefixes
                prefix 2001:db8:ffff::/48 wan-host
            exit
        exit
        group-interface "group-int-1" wlangw create
            ipv6
                allow-multiple-wan-addresses
                router-advertisements
                    managed-configuration
                    other-stateful-configuration
                    dns-options
                        include-dns
                    exit
                    prefix-options
                        autonomous
                        no on-link
                    exit
                    no shutdown
                exit
                router-solicit
                    user-db "ludb-1"
                    no shutdown
                exit
```

```
                dhcp6
                    user-db "ludb-1"
                    proxy-server
                        no shutdown
                    exit
                    override-slaac
                exit
            exit
            ipoe-linking
                shutdown
            exit
            sap-parameters
                sub-sla-mgmt
                    def-sub-id use-auto-id
                    sub-ident-policy "policy-sub-ident-1"
                exit
            exit
            dhcp
                proxy-server
                    emulated-server 172.16.0.1
                    no shutdown
                exit
                lease-populate 10000
                user-db "ludb-1"
                no shutdown
            exit
            ip-mtu 1454
            wlan-gw
                gw-address 172.16.74.244
                gw-ipv6-address 2001:db8::1:1
                router 1
                tcp-mss-adjust 1400
                wlan-gw-group 1
                no shutdown
            exit
        exit
    exit
    no shutdown
exit
```

If the UE is successfully connected, two IPoE hosts will exist on the WLAN-GW:

```
*A:WLAN-GW # show service active-subscribers
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber user-1 (sub-profile-1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[4/2/nat-out-ip:2049.4] - sla:sla-profile-1
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
10.255.0.1
                b0:9f:ba:b9:40:f8 N/A       DHCP
2001:db8:ffff::1/128
```

```
                b0:9f:ba:b9:40:f8 N/A       IPoE-DHCP6
-------------------------------------------------------------------------------
Number of active subscribers : 1
-------------------------------------------------------------------------------
```

The origin of the DHCPv6 lease is shown by issuing:

```
*A:WLAN-GW # show service id 2 dhcp6 lease-state
===============================================================================
DHCP lease state table, service 2
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LeaseTime  Origin   Stdby
-------------------------------------------------------------------------------
2001:db8:ffff::1/128
                b0:9f:ba:b9:40:f8 [4/2/nat-out-ip:20* 23h59m29s  Radius
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

Conclusion

The WLAN-GW supports IPv4/v6 dual-stack UEs. Although the IPv6 support for UEs can handle single-stack IPv6-only UEs, the UEs only have IPv6 support as an add-on to IPv4.

WiFi Aggregation and Offload — Migrant User Support

This chapter provides information about WiFi aggregation and offload for migrant user support configurations.

Topics in this chapter include:

Applicability

This example is applicable to all 7750 SR platforms supporting WLAN gateway (WLAN-GW) IOMs (IOM3-XP and 2xISAs). It provides a functional description of “migrant user” handling on 7750 WLAN-GW and the corresponding configuration. It assumes the user is aware of the general operations and configuration of the basic WLAN-GW function as already described in the 7750 SR OS Triple Play guide.

The configuration with migrant user support enabled was tested on release 11.0R4.

Overview

The term “Migrant user” refers to user equipment (UEs) that connects to a WiFi network service set identification (SSID) but moves 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 long enough to get a DHCP lease from the WLAN-GW. In actual WiFi deployments with portal authentication, it has been observed that a large percentage of users are migrant such that they get a DHCP lease but do not initiate or complete authentication.

Prior to this feature, an Enhanced Subscriber Management (ESM) host is created when the DHCP process completes. This results in the consumption of resources on both the CPM and IOM, limiting the ESM scale and performance for fully authenticated active users. This feature adds support to create an ESM host only after a user has been fully authenticated, either via a web portal or with an AAA server based on completing EAP exchange. In addition, with this feature L2-aware NAPT is required, such that each UE gets the same shared configured inside IP address from the ISA via DHCP. Until a user is authenticated, forwarding of user traffic is constrained (via policy) to DNS and portal server access only.

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 Change of Authorization (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 the RADIUS server remembers the MAC address of the UE from a previous successful portal authentication) then the initial access-accept from RADIUS will trigger the creation of the ESM host.

Migrant User Support for Open SSID Based on Portal Authentication

Sequence Of Events

1. DHCP Is Received From UE On ISA

   Based on the DHCP and L2-aware NAT configuration on the ISA, an IP address is assigned to the user via DHCP. The DHCP and L2-aware NAT configuration is under the soft-gre node under the group-interface, or under vlan-tag range under the soft-gre node on the group-interface.

   A different DHCP lease-time can be configured for an un-authenticated user (initial-lease-time) and an authenticated user (active-lease-time) for which an ESM host has been created. It is suggested that the initial lease be configured to a smaller value while the UE is migrant so that resources can be reclaimed quickly for a truly migrant user that will not complete authentication.

In addition to lease-times, DHCP return options, for example primary and secondary DNS and NBNS server addresses, that can be configured. This configuration can be per soft-GRE group interface or per VLAN range (where a VLAN tag corresponds to an SSID).

Up to 512 bytes of received DHCP options from clients are stored on the ISA. Once the DHCP ACK is sent back to the UE from the ISA, the UE will be created on the ISA in “migrant (or unauthenticated) state”.

A configured L2-aware IP address is returned to each UE and a temporary L2-aware host is created on the anchor ISA for the UE. The NAT policy applicable to this L2-aware NAT for UE in migrant state is also configured under the group-interface (under soft-gre node or under vlan-tag range).

ARP requests coming from the UE in migrant state will be responded to from the ISA. The authentication to RADIUS is triggered on receiving the first Layer 3 data packet as opposed to on a DHCP DISCOVER.

2. Layer 3 Data Packet Received on the ISA

   The first Layer 3 packet (other than DHCP) will trigger RADIUS authentication from the ISA based on configured **isa-radius-policy** in the **configure>aaa** context. The user-name in the access-request is as per the user-name-format configured in the isa-radius-policy. By default it is the MAC address of the UE. The isa-radius-policy can be configured as the authentication policy under the soft-gre group-interface, or under specific VLAN tag ranges on the soft-gre group-interface. The latter allows for the use of a different authentication policy per SSID.

   The RADIUS packets from the ISA are sourced with the IP address owned by the ISA. Each ISA in the WLAN-GW group gets an IP address from a set of contiguous addresses, the start of which is configurable in isa-radius-policy. The nas-ip-address sent in access-request message is configurable in the isa-radius-policy as the ISA’s local IP address or the system IP address. In case the RADIUS server is behind a load-balancer which updates the source IP address of the RADIUS messages, the RADIUS server may use nas-ip-address to route the RADIUS response back. In this case the nas-ip-address should be configured as the ISA’s IP address otherwise the response would incorrectly be routed to the CPM instead of the ISA.

   The debug output below shows a RADIUS accept-request being sent to the RADIUS server on reception of first Layer 3 packet. The debug can be enabled by issuing:

```
debug router "management" radius packet-type authentication | accounting | coa

253 2013/08/07 20:58:35.53 UTC MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 11830
   Info:     anchor egressing frame
             radius-auth-req

   IP/UDP:   from 192.168.0.2:1142 to 192.0.2.3:1812
```

```
RADIUS:   Access-Request (continued)
"

254 2013/08/07 20:58:35.53 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 192.168.0.2:1142 id 40 len 158 vrid 1
    NAS IP ADDRESS [4] 4 192.0.2.3
    NAS PORT TYPE [61] 4 Virtual(5)
    NAS PORT ID [87] 43 GRE rtr-3#lip-192.168.0.1#rip-192.0.2.1
    USER NAME [1] 17 00:0a:0a:00:01:00
    PASSWORD [2] 16 rCmhFboYeM2M8hOuBYJXJk
    CALLING STATION ID [31] 17 00:0a:0a:00:01:00
    VSA [26] 19 Alcatel(6527)
      CHADDR [27] 17 00:0a:0a:00:01:00
"
```

Received Layer 3 packets from the UE are handled as per the redirect-policy configured under the soft-gre group-interface or under applicable VLAN tag range on the soft-gre interface.

The redirect-policy is an IP ACL that should contain one more “forward rules” for traffic that should be forwarded while the UE is pending portal authentication. This typically should include traffic to and from DNS and web portal and is subjected to temporary L2-aware NAT. The redirect-policy also specifies the URL for redirecting triggered by http packets. The redirect-policy and/or the redirect URL can also be overridden via the RADIUS access-accept. Any other non-http traffic that does not match the forward rules is dropped.

While a UE is pending portal authentication no accounting messages are sent to the AAA server. Disconnect-Message from AAA server is supported while the UE is pending authentication.

3. Access-accept from RADIUS

   The access-accept is received on the ISA from which the access-request was generated. The initial access-accept from RADIUS can indicate if a user needs to be authenticated by the portal or is a pre-authenticated user. The indication is based on inclusion of a “redirect policy” applicable to the user in a vendor specific attribute (VSA) (Alc-Wlan-Portal-Redirect, type = string) received from the RADIUS server. 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 the redirect policy that is locally specified under the soft-gre interface or under vlan-tag ranges on soft-gre interface. The redirect-policy is created under sub-mgmt node.

   The UE state is changed to “portal” to indicate the UE is pending portal authentication and has limited access.

   The debug below shows the RADIUS accept-request being received from the RADIUS server and being processed by the WLAN-GW.

```
255 2013/08/07 20:58:35.61 UTC MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 2/1, SeqNo 11831
   Info:     anchor ingressing frame
             portal auth-accept

   IP/UDP:   from 192.0.2.3:1812 to 192.168.0.2:1142

RADIUS:   Access-Accept (continued)
"

256 2013/08/07 20:58:35.62 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 40 len 64 from 192.0.2.3:1812 vrid 1
    VSA [26] 14 Alcatel(6527)
      SUBSC ID STR [11] 12 migrant_user
    VSA [26] 18 Alcatel(6527)
      WLAN PORTAL REDIRECT [172] 16 redirect-policy-1
"
```

The following command is used to display UE information on the ISA, including the state of the UE and the GRE tunnel to the AP through which the UE is connected.

```
*A:PE-1# tools dump wlan-gw ue
===============================================================================
Matched 1 session on Slot #2 MDA #1
===============================================================================
UE-Mac          : 00:0a:0a:00:01:00     UE-vlan         : N/A
UE IP Addr      : 10.0.0.10             Description     : Portal
UE timeout      : 288 sec               Auth-time       : 08/07/13 20:58:35
Tunnel MDA      : 2/2                   Tunnel Router   : 10
MPLS label      : 3000                  Shaper          : Default
GRE Src IP Addr : 192.0.2.2             GRE Dst IP Addr : 192.168.0.1
Anchor SAP      : 2/1/nat-out-ip:2049.1
Last-forward    : None                  Last-move       : None
Rx Frames       : 0                     Rx Octets       : 0
Tx Frames       : 0                     Tx Octets       : 0
-------------------------------------------------------------------------------
===============================================================================
No sessions on Slot #2 MDA #2 match the query
```

If neither of the two redirect related VSAs are included in access-accept, then this indicates a “pre-authenticated user”, and an ESM host is created for the subscriber with a subscriber-profile and other subscriber configuration from access-accept; from here normal ESM based forwarding occurs for the subscriber.

If a user is determined as a “pre-authenticated user”, a message is generated to the CPM to create an ESM host. The information received from RADIUS in the access-accept message (for example subscriber-profile, app-profile etc) and the information from DHCP (for example the DHCP options) are passed in this message.

- COA from RADIUS

  When user's credentials entered on the portal are successfully verified, the portal triggers the AAA server to generate COA to WLAN-GW. The COA serves as a trigger to create an ESM host. The COA MUST contain the subscriber-id and user-name, which are used as a key to identify the UE pending portal authentication.

  The following shows an example debug of a COA being received from the AAA server.

```
248 2013/08/07 19:12:38.29 UTC MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Change of Authorization(43) 192.0.2.3:36776 id 124 len 96 vrid 1
    VSA [26] 19 Alcatel(6527)
      SUBSC ID STR [11] 17 00:0a:0a:00:01:00
    USER NAME [1] 17 00:0a:0a:00:01:00
    VSA [26] 10 Alcatel(6527)
      SLA PROF STR [13] 8 sla-profile-1
    VSA [26] 10 Alcatel(6527)
      SUBSC PROF STR [12] 8 sub-profile-1
"
```

  When the COA is received and successfully processed, a COA-ACK is sent back to the AAA server. The COA message is passed to the CPM to create an ESM host. The information received in the COA, as well as stored information from DHCP (for example the DHCP options) are passed in this message. Once the ESM host is successfully created, the state of the UE on the ISA is changed accordingly to "ESM-user", and can be seen in the output of **tools dump WLAN-GW UE** command, as shown below.

  The UE now has full access (and is not restricted by the original redirect-policy). The COA provides a reference to a subscriber profile that contains the NAT policy for an authenticated UE. The UE continues to keep the same inside L2-aware IP address that was provided originally via DHCP on the ISA. However, the NAT for an authenticated user could be an L2-aware 1:1 NAT or NAPT with a different outside pool and outside ports than the UE in migrant state. The ESM host that is created as described above will also result in the creation of a normal L2-aware host. The original temporary L2-aware host is retained for 10 seconds (and then deleted) to ensure the http response from the portal can be successfully routed back to the UE on the existing connection.

```
 A:PE-1# tools dump wlan-gw ue
===============================================================================
Matched 1 session on Slot #2 MDA #1
===============================================================================
UE-Mac          : 00:0a:0a:00:01:00     UE-vlan         : N/A
UE IP Addr      : N/A                   Description     : ESM-user
UE timeout      : N/A                   Auth-time       : 08/07/13 19:12:38
Tunnel MDA      : 2/2                   Tunnel Router   : 10
MPLS label      : 3000                  Shaper          : 1
GRE Src IP Addr : 192.0.2.2             GRE Dst IP Addr : 192.168.0.1
```

```
Anchor SAP      : 2/1/nat-out-ip:2049.1
Last-forward    : 08/07/13 19:12:25     Last-move       : None
Rx Frames       : 1                     Rx Octets       : 88
Tx Frames       : 1                     Tx Octets       : 222
-------------------------------------------------------------------------------
===============================================================================
No sessions on Slot #2 MDA #2 match the query
```

If UE goes out of range such that the idle timeout expires, the ESM host is deleted and an accounting-stop is sent to the AAA server. If a UE then comes back, and still has a valid DHCP lease, it may not send DHCP DISCOVER or REQUEST and continue to send data. The **data-triggered-ue-creation** command can be configured under soft-gre node on the group-interface (or under vlan-tag ranges on the group-interface) to trigger authentication and recreation of the ESM host for this UE.

The overall sequence of events to take a UE from migrant to authenticated state, where the forwarding of UE traffic is not restricted, is shown in Figure 235.

Figure 235 **Sequence of Events to Establish and Authenticate a Migrant User (continued)**

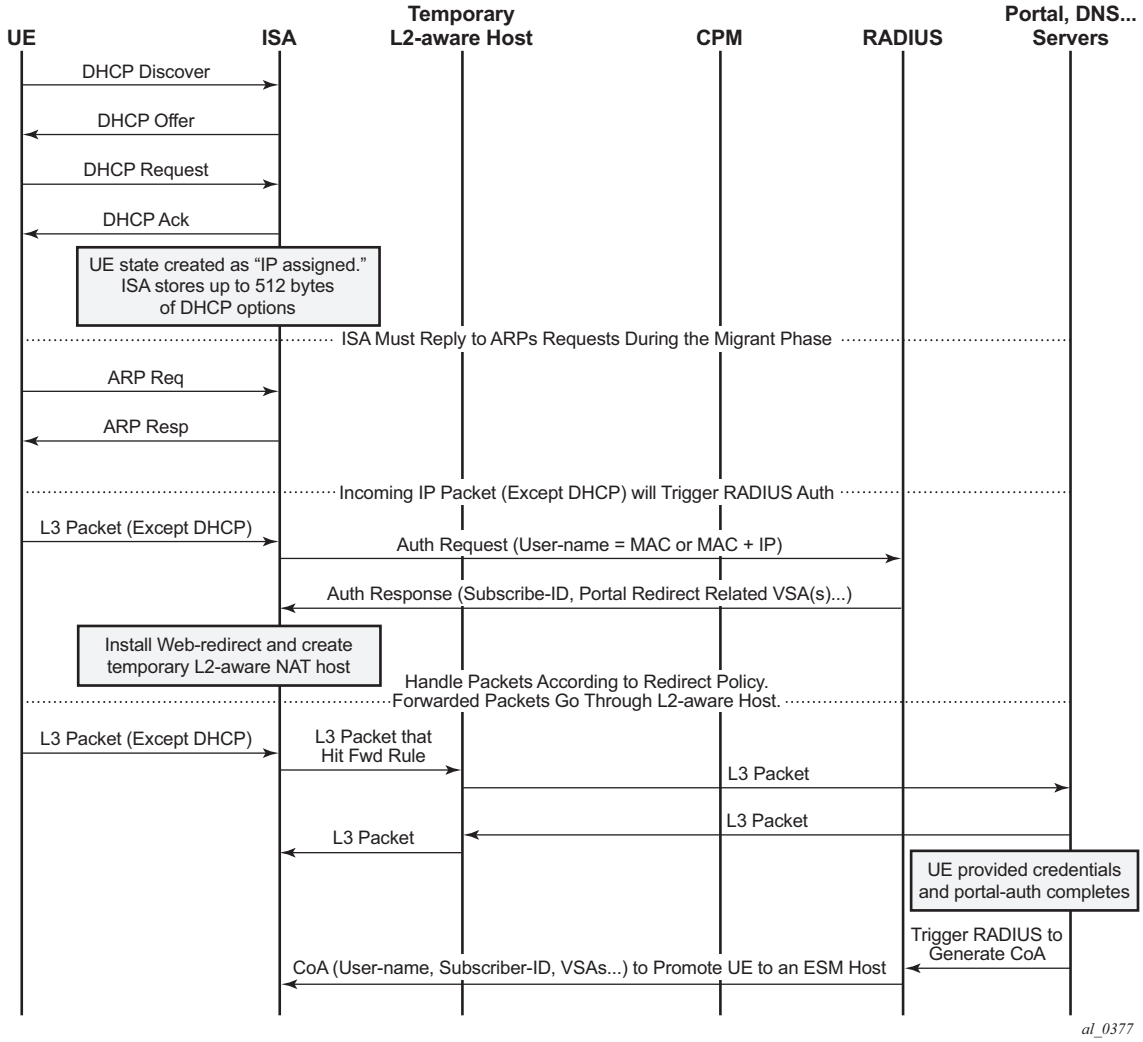

Figure 236 **Sequence of Events to Establish and Authenticate a Migrant Use**

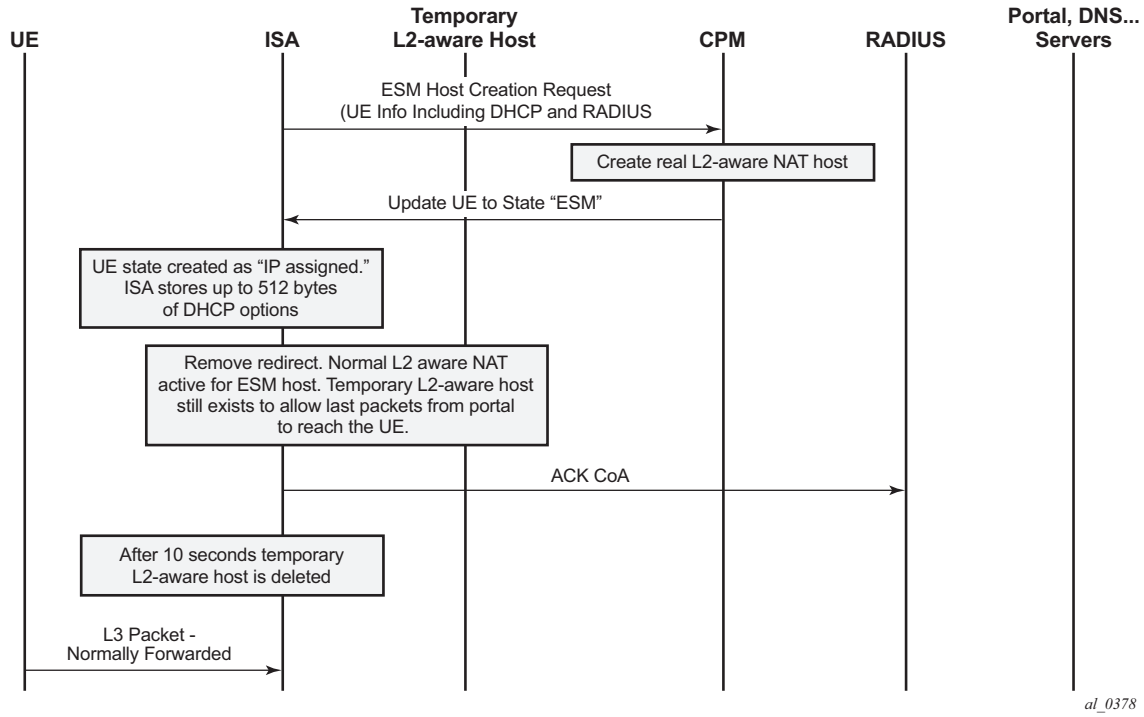

Configuration

The authentication-policy as shown below is used to configure a RADIUS server, and is applicable to the UEs in authenticated state.

```
subscriber-mgmt
   authentication-policy "authentication-1" create
    password "E40PedK6aqrEIpr2DEoJyVR8PQ3XkFF7" hash2
      radius-authentication-server
               source-address 192.0.2.1
               router "management"
               server 1 address 192.0.2.3 secret "6uuGli25Vtl49q0." hash2
      exit
      accept-authorization-change
      include-radius-attribute
               acct-session-id
               circuit-id
               remote-id
               nas-port-id
               nas-identifier
               nas-port-type
               pppoe-service-name
               dhcp-options
               dhcp-vendor-class-id
```

```
            access-loop-options
            mac-address
            called-station-id
            calling-station-id sap-string
            tunnel-server-attrs
```

An isa-radius-policy is required for authentication from the ISA, as below – this contains the attributes to be sent in the access request message to the RADIUS server, which is also configured in this policy.

```
aaa
   isa-radius-policy "isa-policy-1" create
       nas-ip-address-origin isa-ip
       password "CAO6ALDnhyBJERE4xnXoW15MQ/hu74x5nDE7F.OJxHM" hash2
       auth-include-attributes
            called-station-id
            calling-station-id
            circuit-id
            dhcp-options
            dhcp-vendor-class-id
            mac-address
            nas-identifier
            nas-port-id
            nas-port-type
            remote-id
       exit
       servers
            router 1
            source-address-range 192.168.0.2
            server 1 create
                authentication
                coa
                ip-address 192.0.2.3
                secret "CAO6ALDnhyBJERE4xnXoW15MQ/hu74x5nDE7F.OJxHM" hash2
                no shutdown
            exit
       exit
  exit
exit
```

The HTTP redirect policy is shown below, this is enforced on ISA while a UE is migrant and contains the configurations defining the forwarding of traffic in this state.

```
subscriber-mgmt
    http-redirect-policy "redirect-policy-1" create
 url "http://66.185.84.163"
 forward-entries
     dst-ip 192.168.1.1 protocol udp dst-port 53
     dst-ip 192.168.1.2 protocol udp dst-port 53
     dst-ip 66.185.84.163 protocol tcp dst-port 80
     dst-ip 10.0.0.1 protocol udp dst-port 67
     dst-ip 10.0.0.1 protocol udp dst-port 68
  exit
     exit
exit
```

The NAT pool configuration for migrant and authenticated UEs is shown below.

```
vprn 10 customer 1 create
   nat
      inside
          l2-aware
              address 10.0.0.1/24
          exit
      exit
      outside
           pool "migrant-pool-1" nat-group 1 type wlan-gw-anchor create
                 address-range 192.168.2.0 192.168.2.255 create
                 exit
                 no shutdown
           exit
           pool "auth-pool-1" nat-group 1 type l2-aware create
                address-range 192.168.3.0 192.168.3.255 create
                exit
                no shutdown
           exit
       exit
   exit
exit
```

The NAT policy for migrant UEs is as follows.

```
service
   nat
        nat-policy "migrant-policy" create
           pool "migrant-pool-1" router 1
           timeouts
              tcp-established min 1
           exit
        exit
    exit
exit
```

Below is the NAT policy for authenticated UEs.

```
service
    nat
         nat-policy "nat-auth-policy-1" create
            pool "auth-pool-1" router 10
         exit
    exit
exit
```

The migrant user configuration under the soft-gre group-interface within the VPRN service is shown below. This includes configuration for authentication, DHCP, and forwarding from the ISA, as defined in the sections above. The migrant user related configuration can be specified per VLAN tag (or range) under soft-gre interface, where each VLAN tag represents an SSID.

```
vprn 1 customer 1 create
   subscriber-interface "sub-int-1" create
         address 10.0.0.1/24
         group-interface "soft-gre-1" softgre create
             sap-parameters
                  sub-sla-mgmt
                       def-sla-profile "sla-profile-1"
                       def-sub-id use-auto-id
                       def-sub-profile "sub-profile-1"
                       sub-ident-policy "sub_ident"
                   exit
              exit
              dhcp
                  proxy-server
                       emulated-server 10.0.0.1
                       lease-time hrs 1
                       no shutdown
                  exit
                  trusted
                  lease-populate 32767
                  gi-address 10.0.0.1
                  no shutdown
              exit

              authentication-policy "authentication-1"
              host-connectivity-verify

              soft-gre
                   authentication
                       authentication-policy "isa-policy-1"
                   exit
                   gw-address 192.168.0.1
                   mobility
                       hold-time 0
                       trigger data iapp
                   exit
                   router 1
                   wlan-gw-group 1
                   vlan-tag-ranges
                       range start 100 end 100
                           authentication
                                authentication-policy "isa-policy-1"
                           exit
                           data-triggered-ue-creation
                           dhcp
                                active-lease-time min 12
                                initial-lease-time min 5
                                l2-aware-ip-address 10.0.0.10
                                primary-dns 192.168.1.1
                                secondary-dns 192.168.1.2
                                no shutdown
                           exit
                           http-redirect-policy "redirect-policy-1"
                           nat-policy "migrant-policy"
                      exit

                  exit
                  no shutdown
              exit
```

```
        exit
    exit
exit
```

Conclusion

Migrant user support is a useful feature that optimizes system resources (public IP addresses, ESM hosts, CPU processing, etc.) to provide the scale and performance required in live hot-spot and home-spot WiFi deployments at peak times.

WiFi Aggregation and Offload — Open SSID with DSM and Lawful Intercept

This chapter provides information about WiFi Aggregation and Offload — Open SSID with DSM and Lawful Intercept.

Topics in this chapter include:

Applicability

The configuration in this example is applicable to all 7750 SR platforms that function as Wireless Local Area Network gateway (WLAN-GW). This configuration makes use of soft-Generic Routing Encapsulation (soft-GRE) as the access technology, which requires one or more WLAN-Input/Output Module IOMs (IOM3-XP and 2 x MS-ISAs), and was tested using SR OS release 12.0.R4.

Summary

WiFi Aggregation and Offload functionality for the 7750 SR has been supported in SR OS 10.0.R1 and later. This includes a RADIUS proxy server with RADIUS proxy cache and support for soft-GRE tunnels.

Initially, WLAN-GW subscribers were implemented using Enhanced Subscriber Management (ESM) on the Control Processing Module (CPM). To achieve higher scalability, subscribers can be implemented using Distributed Subscriber Management (DSM) on the Multi-Service Integrated Service Adapter (MS-ISA) cards, as described in this chapter.

Law enforcement agencies often require operators to provide a method of intercepting traffic from specific User Equipment (UE). This chapter describes a method of configuring Lawful Intercept (LI) for a DSM UE.

Overview

Starting with release 12.0.R4, DSM can be used for higher scalability by instantiating subscribers on the MS-ISA cards, even after authentication, instead of creating them on the CPM as when using ESM. Therefore, the maximum number of UEs per WLAN-GW, and other performance factors such as setup rate, are higher. When using DSM, commands that are different from those used with ESM are used to monitor the UEs. These commands are similar to those used by the previously available migrant users feature, which only instantiated the users on the MS-ISA cards prior to their authentication.

Lawful intercept can also be configured to intercept traffic to and from a UE. For security reasons, the configuration for LI can be kept separate and invisible to the regular admin user of the 7750 SR, even if this admin user has full admin access. In this situation, any information about the intercept is only available to the LI user. An example LI configuration is shown in this document, with the intercept configured using CLI. It is also possible to trigger DSM LI using RADIUS (in Access-Accept or Change of Authorization (CoA) messages). The RADIUS attributes are encrypted so that knowledge of the intercept cannot be gained by examining a packet capture.

Configuration

The WiFi offload scenario with open SSID, DSM and LI shown in Figure 237 has following characteristics:

- Open SSID with web portal authentication
- DSM with fixed IPv4 address for all UEs, with L2-aware NAT
- Access Point (AP) access in Virtual Private Routed Network (VPRN) 1000
- UEs terminated in VPRN 2004
- Lawful Intercept with separate user account for monitoring or configuring intercepts using the CLI

Figure 237 **WiFi Offload Scenario with Open SSID, DSM and LI**

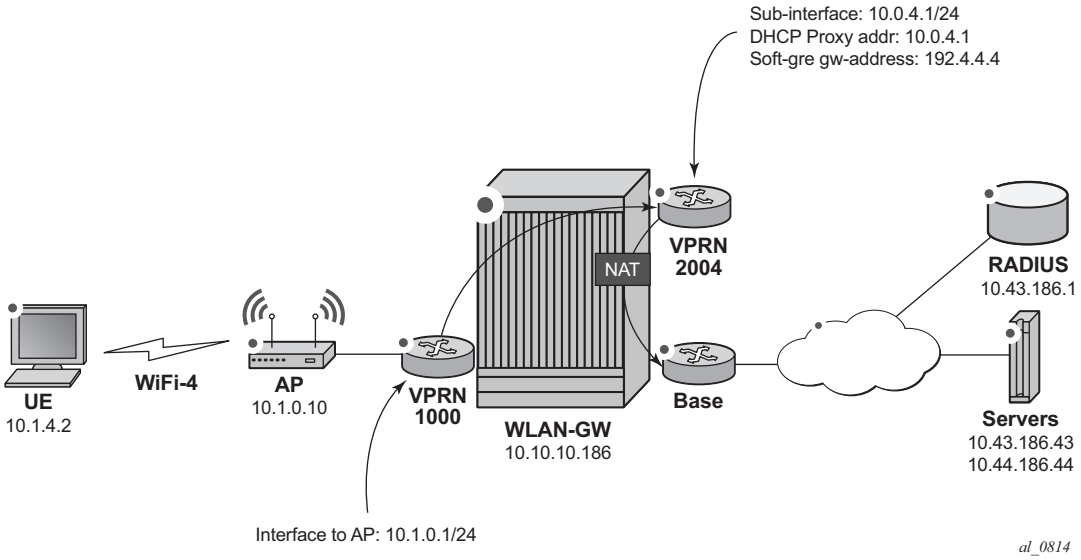

WLAN-GW

Note that the uplink interface, Interior Gateway Protocol (IGP), and system configuration is outside the scope of this document.

The following Card and Media Dependent Adapter (MDA) configuration only shows the WLAN-Input/Output Module (IOM). An IOM3-XP containing two MS-ISA cards provides the WLAN-GW functionality. The MDA type for the ISA cards is isa-bb, the same type that is used for NAT.

```
*A:WLAN-GW# /configure card 3
*A:WLAN-GW>config>card# info
----------------------------------------------
        card-type iom3-xp
        mda 1
            mda-type isa-bb
            no shutdown
        exit
        mda 2
            mda-type isa-bb
            no shutdown
        exit
        no shutdown
----------------------------------------------
```

The following ISA configuration applies.

```
*A:WLAN-GW# /configure isa
*A:WLAN-GW>config>isa# info
----------------------------------------------
        wlan-gw-group 1 create
            active-iom-limit 1
            iom 3
            no shutdown
        exit
----------------------------------------------
```

The following NAT configuration provides an outside pool of type wlan-gw-anchor required to support DSM.

```
*A:WLAN-GW# /configure router nat
*A:WLAN-GW>config>router>nat# info
----------------------------------------------
            outside
                pool "WiFi-4-dsm" nat-group 1 type wlan-gw-anchor create
                    port-reservation ports 15
                    address-range 10.0.40.0 10.0.40.255 create
                    exit
                    no shutdown
                exit
            exit
----------------------------------------------

*A:WLAN-GW# /configure service nat
*A:WLAN-GW>config>service>nat# info
----------------------------------------------
            nat-policy "WiFi-4-dsm" create
                pool "WiFi-4-dsm" router Base
            exit
----------------------------------------------
```

The following Authentication, Authorization and Accounting (AAA) configuration contains an isa-radius-policy. The nas-ip-address-origin parameter selects the IP address sent as the Network Access Server (NAS) IP Address attribute in the ISA RADIUS requests. The source-address-range configures the IP address used by the first MS-ISA card on this WLAN-GW to send and receive RADIUS messages. Other WLAN IOM MS-ISA cards will get consecutive IP addresses in order of slot number. The password and secret configured here have to match the RADIUS server configuration.

```
*A:WLAN-GW# /configure aaa
*A:WLAN-GW>config>aaa# info
----------------------------------------------
        isa-radius-policy "IRS_4" create
            nas-ip-address-origin isa-ip
            password "7USmr6f7JkxD5zb3MeEZnjf1BSqaZkcH" hash2
            servers
                access-algorithm round-robin
                router "Base"
                source-address-range 10.10.186.1
```

```
                server 1 create
                    authentication
                    coa
                    ip-address 10.43.186.1
                    secret "7USmr6f7JkxD5zb3MeEZnjf1BSqaZkcH" hash2
                    no shutdown
                exit
            exit
        exit
-----------------------------------------------
```

The following subscriber management configuration contains the http-redirect-policy for redirecting newly connected UEs to the web portal, and allowing only traffic to the web portal IP address and the Domain Name Server (DNS) server.

```
*A:WLAN-GW# /configure subscriber-mgmt
*A:WLAN-GW>config>subscr-mgmt# info
-----------------------------------------------
        http-redirect-policy "WiFi-4-dsm-redirect" create
            url "http://portal1.3ls.net/portal4.php?mac=$MAC"
            forward-entries
                dst-ip 10.43.186.1 protocol tcp dst-port 80 prefix-length 32
                dst-ip 10.43.186.43 protocol udp dst-port 53 prefix-length 32
            exit
        exit
-----------------------------------------------
```

The following policy configuration is used for exporting required routes, including the address of the APs, the outside NAT prefixes, and the ISA RADIUS source IP addresses.

```
*A:WLAN-GW# /configure router policy-options
*A:WLAN-GW>config>router>policy-options# info
-----------------------------------------------
            prefix-list "WiFi"
                prefix 10.0.0.0/16 longer
                prefix 10.10.186.0/24 longer
            exit
            prefix-list "WiFi-APs"
                prefix 10.1.0.0/16 longer
            exit
            policy-statement "toisis"
                entry 10
                    from
                        prefix-list "WiFi" "WiFi-APs"
                    exit
                    action accept
                    exit
                exit
            exit
            policy-statement "WiFi-APs"
                entry 10
                    from
                        prefix-list "WiFi-APs"
                    exit
                    action accept
```

```
                    exit
                exit
            exit
-----------------------------------------------
```

The following configuration is used for exporting the outside NAT prefixes and the ISA RADIUS source IP addresses to ISIS.

```
*A:WLAN-GW# /configure router isis
*A:WLAN-GW>config>router>isis# info
-----------------------------------------------
        export "toisis"
-----------------------------------------------
```

The following configures VPRN 1000 for AP connectivity, where the AP prefix is exported to the Global Route Table (GRT) so that it can be managed from servers reachable through the Base router.

```
*A:WLAN-GW# /configure service vprn 1000
*A:WLAN-GW>config>service>vprn# info
-----------------------------------------------
            route-distinguisher 65400:1000
            interface "toAP" create
                address 10.1.0.1/24
                sap 1/1/7 create
                exit
            exit
            grt-lookup
                enable-grt
                    static-route 0.0.0.0/0 grt
                exit
                export-grt "WiFi-APs"
            exit
-----------------------------------------------
```

The following configures VPRN 2004 for UE termination, with distributed-sub-mgmt enabled and the ISA RADIUS policy configured under vlan-tag-ranges.

```
*A:WLAN-GW# /configure service vprn 2004
*A:WLAN-GW>config>service>vprn# info
-----------------------------------------------
            description "Open WiFi with DSM"
            route-distinguisher 65400:2004
            subscriber-interface "SI4" create
                address 10.0.4.1/24
                group-interface "GI4" wlangw create
                    wlan-gw
                        gw-address 192.4.4.4
                        mobility
                            trigger data iapp
                        exit
                        router 1000
                        wlan-gw-group 1
                        vlan-tag-ranges
                            range default
```

```
                                    authentication
                                        authentication-policy "IRS_4"
                                    exit
                                    dhcp
                                        active-lease-time min 5
                                        initial-lease-time min 5
                                        l2-aware-ip-address 10.1.4.2
                                        primary-dns 10.43.186.43
                                        secondary-dns 10.44.186.44
                                        no shutdown
                                    exit
                                    distributed-sub-mgmt
                                        no shutdown
                                    exit
                                    nat-policy "WiFi-4-dsm"
                                exit
                            exit
                            no shutdown
                        exit
                    exit
                exit
                nat
                    inside
                        l2-aware
                            address 10.1.4.1/24
                        exit
                    exit
                exit
                wlan-gw
                exit
                no shutdown
----------------------------------------------
```

The following LI user configuration allows user LIadmin to configure and view the Lawful Intercept configuration.

```
*A:WLAN-GW# /configure system security
*A:WLAN-GW>config>system>security# info
----------------------------------------------
            profile "li"
                default-action deny-all
                li
                entry 1
                    match "back"
                    action permit
                exit
                entry 2
                exit
                entry 10
                    match "configure system security"
                    action permit
                exit
                entry 20
                    match "configure li"
                    action permit
                exit
                entry 30
                    match "show li"
```

```
                    action permit
                exit
                entry 40
                    match "file"
                    action permit
                exit
                entry 50
                    match "info"
                    action permit
                exit
                entry 60
                    match "admin display-config"
                    action permit
                exit
                entry 70
                    match "tools perform security"
                    action permit
                exit
                entry 80
                    match "tools dump li wlan-gw ue"
                    action permit
                exit
                entry 100
                    match "exit"
                    action permit
                exit
            exit
            user "LIadmin"
                password "$2y$10$Yp3sQZpGlbg6K3CeQoCHi.wyBOj7ts5/tsY/
nqb0bbFjuFZ9G5wsi"
                access console li
                console
                    no member "default"
                    member "li"
                exit
            exit
-----------------------------------------------
```

The following mirror configuration of type ip-only forwards intercepted traffic to a server.

```
*A:WLAN-GW# /configure mirror
*A:WLAN-GW>config>mirror# info
-----------------------------------------------
        mirror-dest 199 type ip-only create
            encap
                layer-3-encap ip-udp-shim create
                    gateway create
                        ip src 10.10.10.186 dest 10.43.186.43
                        udp src 3199 dest 3199
                    exit
                exit
            exit
            no shutdown
        exit
-----------------------------------------------
```

The following configures a BOF, with li-local-save, a local LI config file, and li-separate, ensuring that only the LI user can view or modify LI parameters.

```
*A:WLAN-GW# show bof
===============================================================================
BOF (Memory)
===============================================================================
    li-local-save
    li-separate
```

The following LI source configuration to intercept the UE can only be configured or viewed by user LIadmin. The configuration is saved in cf3:\li.cfg, encrypted. The intercept-id and session-id will appear in the LI packets, which can be decoded in Wireshark using Decode As, Jmirror.

```
*A:WLAN-GW# /configure li
*A:WLAN-GW>config>li# info
----------------------------------------------
#--------------------------------------------------
echo "LI Log Configuration"
#--------------------------------------------------
        log
        exit
#--------------------------------------------------
echo "LI Filter Lock State Configuration"
#--------------------------------------------------
        li-filter-lock-state locked
#--------------------------------------------------
echo "LI Mirror Source Configuration"
#--------------------------------------------------
        li-source 199
            wlan-gw
                dsm-subscriber mac 68:7f:74:8b:3d:d7
                    intercept-id 1
                    session-id 199
                exit
            exit
            no shutdown
        exit
----------------------------------------------
```

Freeradius

This default configuration section sets the VSA Alc-Wlan-Ue-Creation-Type with value 1, which triggers the creation of a DSM host (value 0 is ESM). The Nas-Ip-Address is returned in the Alc-Wlan-Portal-URL to tell the web portal which MS-ISA address should receive the RADIUS CoA request:

```
/etc/freeradius/users:

DEFAULT         Auth-Type := Local, User-Password := "alcatel", user-name=~"^.*$"
                        Alc-Subsc-ID-Str = "%{User-Name}",
```

```
                        Alc-Wlan-Ue-Creation-Type = 1,
                        Alc-Wlan-Portal-Redirect = "WiFi-4-dsm-redirect",
                        Alc-Wlan-Portal-URL = "http://portal1.3ls.net/
portal4.php?nas=%{Nas-Ip-Address}&mac=%{User-Name}&ssid=WiFi-4",
```

As an alternative to configuring the LI for the UE in the CLI, the following RADIUS attributes can be sent in the Access-Accept and CoA.

```
                        Alc-LI-Action = "enable",
                        Alc-LI-Destination = "199",
                        Alc-LI-Intercept-Id = 1,
```

In /etc/freeradius/clients.conf, each ISA is a client.

```
client 10.10.186.1 {
        secret      = alcatel
        shortname   = WLAN-GW-ISA1
}
client 10.10.186.2 {
        secret      = alcatel
        shortname   = WLAN-GW-ISA2
}
```

A RADIUS CoA sent during a successful portal login makes the following UE a DSM subscriber with full access.

```
echo "User-Name='.$mac.',Alc-Wlan-Ue-Creation-Type="1",Alc-Subsc-Prof-
Str="SUBP_4",Alc-SLA-Prof-Str=SLAP_4,Alc-Primary-Dns = 10.43.186.43," | /usr/bin/
radclient -x -r 1 -t 2 '.$nas.' coa alcatel
```

Access Points

The following must be configured on the Access Point as a minimum:

- IP address 10.1.1.10/24
- Default route to 10.1.1.1
- Open SSID WiFi-4 mapped to VLAN 40
- Soft-GRE tunnel with destination 192.4.4.4, with VLAN 40 mapped to this tunnel

Show Commands

The following commands show the status of the UEs. For DSM users, the UEs are displayed using a tools command. Before portal authentication, the UE is in Portal state.

```
*A:WLAN-GW# /tools dump wlan-gw ue
===============================================================================
Matched 1 session on Slot #3 MDA #1
===============================================================================
UE-Mac          : 68:7f:74:8b:3d:d7     UE-vlan         : 40
UE IP Addr      : 10.1.4.2              UE timeout      : 293 sec
Description     : Portal
Auth-time       : 09/08/2014 11:30:34
Tunnel MDA      : 3/2                   Tunnel Router   : 1000
MPLS label      : 40                    Shaper          : Default
Tunnel Src IP   : 10.1.0.10             Tunnel Dst IP   : 192.4.4.4
Tunnel Type     : GRE
Anchor SAP      : 3/1/nat-out-ip:2049.3
AP-Mac          : Unknown               AP-RSSI         : Unknown
AP-SSID         : Unknown
Last-forward    : 09/08/2014 11:30:39   Last-move       : None
Session Timeout : None                  Idle Timeout    : N/A
Acct Update     : None                  Acct Interval   : N/A
Acct Session-Id : N/A
Acct Policy     : N/A
NAT Policy      : WiFi-4-dsm
Redirect Policy : WiFi-4-dsm-redirect
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       : 204                   Rx Octets       : 17381
Tx Frames       : 78                    Tx Octets       : 67793
-------------------------------------------------------------------------------
===============================================================================
No sessions on Slot #3 MDA #2 match the query
```

After login to the web portal, the UE transitions to a DSM-user and the Redirect Policy is removed,

```
*A:WLAN-GW# /tools dump wlan-gw ue
===============================================================================
Matched 1 session on Slot #3 MDA #1
===============================================================================
UE-Mac          : 68:7f:74:8b:3d:d7     UE-vlan         : 40
UE IP Addr      : 10.1.4.2              UE timeout      : 284 sec
Description     : DSM-user
Auth-time       : 09/08/2014 11:30:45
Tunnel MDA      : 3/2                   Tunnel Router   : 1000
MPLS label      : 40                    Shaper          : Default
Tunnel Src IP   : 10.1.0.10             Tunnel Dst IP   : 192.4.4.4
Tunnel Type     : GRE
Anchor SAP      : 3/1/nat-out-ip:2049.3
AP-Mac          : Unknown               AP-RSSI         : Unknown
AP-SSID         : Unknown
Last-forward    : 09/08/2014 11:30:47   Last-move       : None
Session Timeout : None                  Idle Timeout    : N/A
Acct Update     : None                  Acct Interval   : N/A
Acct Session-Id : N/A
Acct Policy     : N/A
NAT Policy      : WiFi-4-dsm
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       : 273                   Rx Octets       : 23186
Tx Frames       : 122                   Tx Octets       : 108083
-------------------------------------------------------------------------------
===============================================================================
No sessions on Slot #3 MDA #2 match the query
```

User LIadmin can view the configured intercept,

```
*A:WLAN-GW>config>li# /tools dump li wlan-gw ue
===============================================================================
Matched 1 session on Slot #3 MDA #1
===============================================================================
UE-Mac          : 68:7f:74:8b:3d:d7     Mirror Service  : 199
LI Intercept-Id : 1                     LI Session-Id   : 199
-------------------------------------------------------------------------------
===============================================================================
No sessions on Slot #3 MDA #2 match the query
```

Debug

In this example, the following debug configuration applies.

```
debug
    mirror-source 99
        port 1/1/7 egress ingress
        port 1/1/9 egress ingress
        no shutdown
    exit
    wlan-gw
        group 1
            ue 68:7f:74:8b:3d:d7 packet dhcp radius
        exit
    exit
exit

The debug trace starts with DHCP.

150 2014/09/08 11:30:31.93 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 3/1, SeqNo 3528
   Info:     anchor ingressing frame
             received upstream from tunnel

   Ethernet: from 68:7f:74:8b:3d:d7 to ff:ff:ff:ff:ff:ff (ethertype: 0x0800)

   IP/UDP:   from 0.0.0.0 (port 68) to 255.255.255.255 (port 67)

   DHCP:
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 68:7f:74:8b:3d:d7    xid: 0xca073331
```

```
   DHCP options:
   [53] Message type: Discover
   [61] Client id: (hex) 01 68 7f 74 8b 3d d7
   [12] Host name: W81VM
   [60] Class id: MSFT 5.0
   [55] Param request list: len = 13
             1  Subnet mask
            15  Domain name
             3  Router
             6  Domain name server
            44  NETBIOS name server
            46  NETBIOS type
            47  NETBIOS scope
            31  Router discovery
            33  Static route
           121  Unknown option
           249  Unknown option
           252  Unknown option
            43  Vendor specific
   [255] End
"

151 2014/09/08 11:30:31.93 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 3/2, SeqNo 706
   Info:     tunnel ingressing frame
             received downstream from anchor

   Ethernet: from 00:00:00:02:02:02 to 68:7f:74:8b:3d:d7 (ethertype: 0x0800)

   IP/UDP:   from 10.1.4.1 (port 67) to 10.1.4.2 (port 68)

   DHCP:
   ciaddr: 0.0.0.0           yiaddr: 10.1.4.2
   siaddr: 10.1.4.1          giaddr: 0.0.0.0
   chaddr: 68:7f:74:8b:3d:d7    xid: 0xca073331

   DHCP options:
   [53] Message type: Offer
   [54] DHCP server addr: 10.1.4.1
   [1] Subnet mask: 255.255.255.0
   [3] Router: 10.1.4.1
   [51] Lease time: 300
   [6] Domain name server: length = 8
             10.43.186.43
             10.44.186.44
   [255] End
"

152 2014/09/08 11:30:32.09 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 3/1, SeqNo 3529
   Info:     anchor ingressing frame
             received upstream from tunnel

   Ethernet: from 68:7f:74:8b:3d:d7 to ff:ff:ff:ff:ff:ff (ethertype: 0x0800)

   IP/UDP:   from 0.0.0.0 (port 68) to 255.255.255.255 (port 67)

   DHCP:
   ciaddr: 0.0.0.0           yiaddr: 0.0.0.0
```

```
   siaddr: 0.0.0.0           giaddr: 0.0.0.0
   chaddr: 68:7f:74:8b:3d:d7    xid: 0xca073331

   DHCP options:
   [53] Message type: Request
   [61] Client id: (hex) 01 68 7f 74 8b 3d d7
   [50] Requested IP addr: 10.1.4.2
   [54] DHCP server addr: 10.1.4.1
   [12] Host name: W81VM
   [81] client FQDN: rcode1: 0, rcode2: 0, domain name = (hex) 00 57 38 31 56
   4d
   [60] Class id: MSFT 5.0
   [55] Param request list: len = 13
             1  Subnet mask
            15  Domain name
             3  Router
             6  Domain name server
            44  NETBIOS name server
            46  NETBIOS type
            47  NETBIOS scope
            31  Router discovery
            33  Static route
           121  Unknown option
           249  Unknown option
           252  Unknown option
            43  Vendor specific
   [255] End
"

153 2014/09/08 11:30:32.09 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 3/2, SeqNo 707
   Info:     tunnel ingressing frame
             received downstream from anchor

   Ethernet: from 00:00:00:02:02:02 to 68:7f:74:8b:3d:d7 (ethertype: 0x0800)

   IP/UDP:   from 10.1.4.1 (port 67) to 10.1.4.2 (port 68)

   DHCP:
   ciaddr: 0.0.0.0           yiaddr: 10.1.4.2
   siaddr: 10.1.4.1          giaddr: 0.0.0.0
   chaddr: 68:7f:74:8b:3d:d7    xid: 0xca073331

   DHCP options:
   [53] Message type: Ack
   [54] DHCP server addr: 10.1.4.1
   [1] Subnet mask: 255.255.255.0
   [3] Router: 10.1.4.1
   [51] Lease time: 300
   [58] Renew timeout: 150
   [59] Rebind timeout: 263
   [6] Domain name server: length = 8
             10.43.186.43
             10.44.186.44
   [255] End
"
```

RADIUS authentication is triggered by the first data packet.

```
154 2014/09/08 11:30:34.76 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 3/1, SeqNo 3563
   Info:     anchor egressing frame
             radius-auth-req

   IP/UDP:   from 10.10.186.1 (port 1082) to 10.43.186.1 (port 1812)

   RADIUS:   Access-Request (continued)
"

155 2014/09/08 11:30:34.76 EDT MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Transmit
  Access-Request(1) 10.10.186.1:1082 id 45 len 126 vrid 1
    NAS IP ADDRESS [4] 4 10.10.186.1
    USER NAME [1] 17 68:7f:74:8b:3d:d7
    PASSWORD [2] 16 MvqAtmAOvSeeWgNIGyT/t.
    CALLING STATION ID [31] 17 68:7f:74:8b:3d:d7
    CALLED STATION ID [30] 17 00:00:00:00:00:00
    VSA [26] 19 Alcatel(6527)
      CHADDR [27] 17 68:7f:74:8b:3d:d7
"

156 2014/09/08 11:30:34.76 EDT MINOR: DEBUG #2001 Base WLAN-GW
"WLAN-GW: MDA 3/1, SeqNo 3564
   Info:     anchor ingressing frame
             portal auth-accept

   IP/UDP:   from 10.43.186.1 (port 1812) to 10.10.186.1 (port 1082)

   RADIUS:   Access-Accept (continued)
"

157 2014/09/08 11:30:34.76 EDT MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Access-Accept(2) id 45 len 258 from 10.43.186.1:1812 vrid 1
    VSA [26] 19 Alcatel(6527)
      SUBSC ID STR [11] 17 68:7f:74:8b:3d:d7
    VSA [26] 20 Alcatel(6527)
    VSA [26] 20 Alcatel(6527)
    VSA [26] 20 Alcatel(6527)
    VSA [26] 6 Alcatel(6527)
      WLAN UE CREATION TYPE [184] 4 1
    VSA [26] 25 Alcatel(6527)
      WLAN PORTAL REDIRECT [172] 23 WiFi-4-dsm-redirect
    VSA [26] 86 Alcatel(6527)
      WLAN PORTAL URL [173] 84 http://portal1.3ls.net/portal4.php?nas=10.10.186.
1&mac=68:7f:74:8b:3d:d7&ssid=WiFi-4
"
```

RADIUS returns the redirect policy and portal URL, and the UE is then in Portal state. Next, the user logs in and a CoA is sent by the portal.

```
159 2014/09/08 11:30:45.07 EDT MINOR: DEBUG #2001 Base RADIUS
"RADIUS: Receive
  Change of Authorization(43) id 220 len 91 from 10.43.186.1:53449 vrid 1
    USER NAME [1] 17 68:7f:74:8b:3d:d7
    VSA [26] 6 Alcatel(6527)
      WLAN UE CREATION TYPE [184] 4 1
```

```
    VSA [26] 8 Alcatel(6527)
      SUBSC PROF STR [12] 6 SUBP_4
    VSA [26] 8 Alcatel(6527)
      SLA PROF STR [13] 6 SLAP_4
    VSA [26] 6 Alcatel(6527)
      PRIMARY DNS [9] 4 10.43.186.43
"
```

Finally the UE is in DSM state with unrestricted access.

Conclusion

The 7750 SR WLAN-GW, with Open SSID, can support WiFi Offload users as DSM subscribers instantiated on MS-ISA cards. This allows the support of a greater number of UEs on a single system when a full ESM feature set is not required. DSM UEs, just as ESM UEs, can have their traffic intercepted using LI.

Customer Document and Product Support

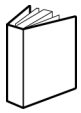

Customer documentation

Customer Documentation Welcome Page

Technical Support

Product Support Portal

Documentation feedback

Customer Documentation Feedback

3HE 11599 AAAA TQZZA 01