



Hewlett Packard
Enterprise

HPE 5940-CMW710-R6710P03 Release Notes

Contents

| | |
|--|----|
| Introduction..... | 1 |
| Version number..... | 1 |
| Version history | 1 |
| Hardware and software compatibility matrix | 32 |
| ISSU compatibility list..... | 34 |
| Upgrading restrictions and guidelines..... | 35 |
| Hardware feature updates | 35 |
| R6710P03 | 35 |
| R6710P01 | 35 |
| R6710..... | 35 |
| F6705..... | 35 |
| E6702..... | 35 |
| R2717..... | 35 |
| R2702..... | 36 |
| F2701 | 36 |
| R2612P02 | 36 |
| R2612P01 | 36 |
| R2612..... | 36 |
| R2610..... | 36 |
| R2609..... | 36 |
| R2608P02 | 36 |
| F2608..... | 36 |
| F2606..... | 36 |
| F2605..... | 36 |
| F2604 | 37 |
| F2603 | 37 |
| R2509P02 | 37 |
| R2509P01 | 37 |
| R2509..... | 37 |
| R2508..... | 37 |
| R2507..... | 37 |
| R2506..... | 37 |
| Software feature and command updates | 37 |
| MIB updates | 38 |
| Operation changes | 41 |
| Operation changes in R6710P03..... | 41 |
| Operation changes in R6710P01..... | 41 |
| Operation changes in R6710 | 41 |
| Operation changes in F6705..... | 42 |
| Operation changes in E6702 | 43 |
| Operation changes in R2717 | 43 |
| Operation changes in R2702 | 43 |
| Operation changes in F2701..... | 44 |
| Operation changes in R2612P02..... | 44 |
| Operation changes in R2612P01..... | 45 |
| Operation changes in R2612 | 46 |
| Operation changes in R2610 | 47 |
| Operation changes in R2609 | 48 |
| Operation changes in R2608P02..... | 48 |
| Operation changes in F2608..... | 50 |
| Operation changes in F2606..... | 52 |
| Operation changes in F2605..... | 52 |
| Operation changes in F2604..... | 53 |

| | |
|---|------------|
| Operation changes in F2603..... | 54 |
| Operation changes in R2509P02..... | 55 |
| Operation changes in R2509P01..... | 55 |
| Operation changes in R2509..... | 56 |
| Operation changes in R2508..... | 56 |
| Operation changes in R2507..... | 56 |
| Operation changes in R2506..... | 56 |
| Restrictions and cautions..... | 56 |
| Restrictions..... | 56 |
| Open problems and workarounds..... | 56 |
| List of resolved problems..... | 57 |
| Resolved problems in R6710P03..... | 57 |
| Resolved problems in R6710P01..... | 62 |
| Resolved problems in R6710..... | 62 |
| Resolved problems in F6705..... | 66 |
| Resolved problems in E6702..... | 69 |
| Resolved problems in R2717..... | 69 |
| Resolved problems in R2702..... | 92 |
| Resolved problems in F2701..... | 99 |
| Resolved problems in R2612P02..... | 99 |
| Resolved problems in R2612P01..... | 100 |
| Resolved problems in R2612..... | 104 |
| Resolved problems in R2610..... | 122 |
| Resolved problems in R2609..... | 126 |
| Resolved problems in R2608P02..... | 133 |
| Resolved problems in F2608..... | 159 |
| Resolved problems in F2606..... | 186 |
| Resolved problems in F2605..... | 187 |
| Resolved problems in F2604..... | 196 |
| Resolved problems in F2603..... | 204 |
| Resolved problems in R2509P02..... | 209 |
| Resolved problems in R2509P01..... | 216 |
| Resolved problems in R2509..... | 216 |
| Resolved problems in R2508..... | 222 |
| Resolved problems in R2507..... | 229 |
| Resolved problems in R2506..... | 232 |
| Related documentation..... | 232 |
| Support and other resources..... | 232 |
| Accessing Hewlett Packard Enterprise Support..... | 232 |
| Documents..... | 233 |
| Documentation feedback..... | 233 |
| Appendix A Feature list..... | 234 |
| Hardware features..... | 234 |
| Software features..... | 237 |
| Appendix B Fixed security vulnerabilities..... | 241 |
| Fixed security vulnerabilities in R6710..... | 241 |
| Fixed security vulnerabilities in F6705..... | 241 |
| Fixed security vulnerabilities in R2702..... | 241 |
| Fixed security vulnerabilities in R2612..... | 242 |
| Fixed security vulnerabilities in R2610..... | 242 |
| Fixed security vulnerabilities in R2609..... | 242 |
| Fixed security vulnerabilities in R2608P02..... | 243 |
| Fixed security vulnerabilities in F2608..... | 243 |
| Fixed security vulnerabilities in F2605..... | 243 |
| Fixed security vulnerabilities in F2604..... | 244 |

| | |
|---|------------|
| Fixed security vulnerabilities in F2603 | 244 |
| Fixed security vulnerabilities in R2509P02 | 245 |
| Fixed security vulnerabilities in R2509 | 245 |
| Fixed security vulnerabilities in R2508 | 245 |
| Fixed security vulnerabilities in R2507 | 246 |
| Appendix C Upgrading software | 247 |
| System software file types | 247 |
| System startup process..... | 248 |
| Upgrade methods..... | 248 |
| Upgrading from the CLI..... | 249 |
| Preparing for the upgrade | 249 |
| Downloading software to the master switch | 250 |
| Upgrading the software images | 253 |
| Installing a patch package..... | 254 |
| Upgrading from the Boot menu | 255 |
| Prerequisites | 255 |
| Accessing the Boot menu | 256 |
| Accessing the basic Boot menu | 257 |
| Accessing the extended Boot menu..... | 258 |
| Using TFTP to upgrade software images through the management Ethernet port..... | 259 |
| Using FTP to upgrade software through the management Ethernet port | 261 |
| Using XMODEM to upgrade software through the console port | 263 |
| Using TFTP to upgrade Boot ROM through the management Ethernet port | 267 |
| Using FTP to upgrade Boot ROM through the management Ethernet port | 268 |
| Using XMODEM to upgrade Boot ROM through the console port | 270 |
| Managing files from the Boot menu | 274 |
| Displaying all files..... | 274 |
| Deleting files..... | 275 |
| Changing the attribute of software images..... | 275 |
| Handling software upgrade failures..... | 277 |

List of Tables

| | |
|---|-----|
| Table 1 Version history | 1 |
| Table 2 Hardware and software compatibility matrix | 32 |
| Table 3 ISSU compatibility list | 34 |
| Table 4 MIB updates | 38 |
| Table 5 5940 series hardware features | 234 |
| Table 6 FF 5940 2-slot Switch/5940 4-slot Switch series hardware features | 235 |
| Table 7 Software features of the 5940 series | 237 |
| Table 8 Shortcut keys | 256 |
| Table 9 Basic Boot ROM menu options | 257 |
| Table 10 BASIC ASSISTANT menu options | 257 |
| Table 11 Extended Boot ROM menu options | 258 |
| Table 12 EXTENDED ASSISTANT menu options | 259 |
| Table 13 TFTP parameter description | 260 |
| Table 14 FTP parameter description | 262 |
| Table 15 TFTP parameter description | 268 |
| Table 16 FTP parameter description | 269 |

Introduction

This document describes the features, restrictions and guidelines, open problems, and workarounds for version 6710P03. Before you use this version in a live network, back up the configuration and test the version to avoid software upgrade affecting your live network.

Use this document in conjunction with *HPE 5940-CMW710-R6710P03 Release Notes (Software Feature Changes)* and the documents listed in "[Related documentation](#)."

Version number

HPE Comware Software, Version 7.1.070, Release 6710P03

Note: You can see the version number with the **display version** command in any view. Please see Note ①.

Version history



IMPORTANT:

The software feature changes listed in the version history table for each version are not complete. To obtain complete information about all software feature changes in each version, see the *Software Feature Changes* document for this release notes.

Table 1 Version history

| Version number | Last version | Release Date | Release type | Remarks |
|-----------------------|-----------------------|--------------|-----------------|---|
| | | | | Added feature includes: <ul style="list-style-type: none">• New feature: DRNI configuraton• New feature: Generating a log message to display product version numbers before and after a software upgrade |
| 5940-CMW710-R 6710P03 | 5940-CMW710-R67 10P01 | 2023-08-01 | Release version | Modified features. <ul style="list-style-type: none">• Modified feature: Enabling SNMP notification for FC interface status changes• Modified feature: Adding a member to a zone alias• Modified feature: Adding a member to a zone• Modified feature: Displaying parent information for a zone member |
| 5940-CMW710-R 6710P01 | 5940-CMW710-R67 10 | 2023-03-02 | Release version | This version fixed bugs. |
| 5940-CMW710-R 6710 | 5940-CMW710-F67 05 | 2022-12-28 | Release version | Added feature includes: <ul style="list-style-type: none">• New features: Fundamentals features• New features: Layer 2—LAN |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|--|
| | | | | <p>switching features</p> <ul style="list-style-type: none"> • New features: Layer 3—IP services features • New features: Layer 3—IP routing features • New features: IP multicast features • New features: MPLS features • New features: ACL and QoS features • New features: Security features • New features: High availability features • New features: Network management and monitoring features • New features: Telemetry features • New features: OpenFlow features • New features: VXLAN features • New features: EVPN features • New features: Intelligent lossless network features • New features: M-LAG support for DRNI commands • New feature: EVPN instance created in system view <p>Modified features.</p> <ul style="list-style-type: none"> • Modified feature: IPv6 routes with prefixes longer than 64 bits • Modified feature: Match criteria in a traffic class • Modified feature: Associating a traffic behavior with a traffic class • Modified feature: Displaying the running configuration • Modified feature: Displaying the contents of the configuration file for the next system startup • Modified feature: Optimized display of BGP BMP server information • Modified feature: Disabling BGP session establishment with peers and peer groups • Modified feature: Optimizations to VXLAN command output • Modified feature: Sharing of VSI interfaces among VSIs |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|---|
| | | | | <ul style="list-style-type: none"> Modified feature: Enabling L2TP for the specified protocol Modified feature: Creating a service loopback group and specify its service type Modified feature: Displaying LSP information Modified feature: Configuring the signaling protocol for an MPLS TE tunnel interface Modified feature: Configuring the tunnel selection order and set the number of tunnels for load balancing Modified feature: Verifying MPLS LSP connectivity for an IPv4 prefix Modified feature: Tracing MPLS LSPs from the ingress node to the egress node for an IPv4 prefix Modified feature: Creating a local site Modified feature: Enabling link flapping protection on an interface Modified feature: AAA methods in an ISP domain Modified feature: Setting the 802.1X periodic reauthentication timer Modified feature: Setting the periodic MAC reauthentication timer Modified feature: Enabling SNMP notifications for L3VPN Modified feature: Displaying BGP VPNv4 routing information Modified feature: Displaying BGP VPNv6 routing information Modified feature: Creating an SNMPv3 user Modified feature: Displaying local public keys Modified feature: Flow-mirroring traffic to an interface Modified feature: Disabling the device from learning the source MAC addresses of Layer 2 protocol packets Modified feature: Restrictions for an interface to join a Layer |

| Version number | Last version | Release Date | Release type | Remarks |
|-------------------|-------------------|--------------|-----------------|--|
| | | | | 2 aggregation group <ul style="list-style-type: none"> Modified feature: Specifying multiple VLAN IDs for VLAN name TLVs Modified feature: Applying a QoS policy to an interface |
| 5940-CMW710-F6705 | 5940-CMW710-E6702 | 2022-07-25 | Feature version | Added feature includes: <ul style="list-style-type: none"> New feature: FEC mode New feature: Displaying ND entry statistics New feature: User IP address conflict SNMP notifications for ARP Modified feature includes: <ul style="list-style-type: none"> Modified feature: Restrictions for an interface to join a Layer 2 aggregation group Modified feature: Configuring MAC address borrowing Modified feature: Configuring the types of advertisable TLVs on a port Modified feature: Enabling IGMP proxying on an interface Modified feature: DRNI term changes Modified feature: Using OSPF commands on interfaces Modified feature: Configuring a static route Modified feature: Configuring an IPv6 static route |
| 5940-CMW710-E6702 | 5940-CMW710-R2717 | 2022-04-14 | ESS version | Added feature includes: <ul style="list-style-type: none"> New feature: Specifying a security enhanced level New feature: Private VSI Modified feature includes: <ul style="list-style-type: none"> Modified feature: Configuring the global priority trust mode for VXLAN packets Modified feature: Configuring MAC authentication Modified feature: Disabling BGP from flushing all routes to the routing table Modified feature: Support for HMAC-SHA-256 OSPF/OSPFv3 authentication mode Modified feature: Displaying the hash keys used for link aggregation load sharing |
| 5940-CMW710-R | 5940-CMW710-R27 | 2021-09-30 | Release | Added feature includes: |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|--|
| 2717 | 02 | | version | <ul style="list-style-type: none"> • New feature: Interface alarm functions • New feature: Establishing neighbors through the secondary IP address of an interface • Advertising only the global unicast address in the NEXT_HOP attribute <p>There are also modified features.</p> <ul style="list-style-type: none"> • New feature: Interface alarm functions • New feature: Establishing neighbors through the secondary IP address of an interface • New feature: Advertising only the global unicast address in the NEXT_HOP attribute • New feature: Assigning a VXLAN tunnel interface to a VXLAN tunnel group • New feature: Setting VLAN tags for sent protocol packets on an interface • New feature: Configuring the system ID • New feature: IPv6 duplicate detection on duplicate addresses • New feature: BGP route re-origination • New feature: Sending RA messages over VXLAN tunnels • New feature: SAVA • New feature: IPv6 IS-IS link tag • New feature: OSPFv3 link tag • New feature: Displaying a message on enablement of the FTP or Telnet server • New feature: Enabling OSPF to advertise the maximum link cost to neighbors • New feature: Enabling IS-IS to advertise the maximum link cost to neighbors • New feature: Enabling OSPFv3 to advertise the maximum link cost to neighbors • New feature: Enabling BGP to advertise a default route to a peer or peer group • New feature: SAVI • New feature: Setting the |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|--|
| | | | | <p>maximum number of IPv6SG bindings on an interface</p> <ul style="list-style-type: none"> • New feature: Setting the interval at which the SNMP module examines the system configuration for changes • New feature: Displaying SNMP notifications drop records • New feature: Displaying SNMP notifications sending records • New feature: Clearing all records from the SNMP trap buffer • New feature: PWDCTL messages • New feature: LIPC messages • New feature: Applying an SSL server policy to the NETCONF over SOAP over HTTPS service • New feature: Policing ND protocol packet to the control plane • New feature: RA message transmission over VXLAN tunnels between distributed EVPN gateways • New feature: MAC move logging for VSIs • New feature: MAC learning limits for VSIs • New feature: DRNI in EVPN-DCI • New feature: L2VPN statistics collection interval • New feature: Unsolicited NA learning • New feature: IPv6 enhancement for OpenFlow flow entries • New feature: IPv6-matching QoS policy • New feature: Enabling BGP to not flush specific routes to the routing table • New feature: Enabling BGP to reset peer sessions gracefully • New feature: Setting the extra time to wait after the restart timer expires • New feature: Enabling mandatory weak password change • New feature: Configuring ND |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|---|
| | | | | <ul style="list-style-type: none"> packet rate limit New feature: gRPC logging in dial-in mode New feature: gRPC logging in dial-out mode New feature: Creating frame match criteria based on VXLAN IDs for dynamic ACs on the IPL of an EVPN DR system New feature: Removing Layer 3 VPN route targets from BGP EVPN routes for VXLAN mapping New feature: Enabling asymmetric IRB mode for EVPN VXLAN traffic forwarding New feature: Setting the Gateway IP Address field for EVPN IP prefix advertisement routes New feature: Enabling BGP to redistribute all IGP routes with the same destination address New feature: Policing ND protocol packets New feature: Using an ACL to control TCP connections from IPv4 HTTP or HTTPS clients New feature: Enabling Layer 3 packet statistics counting New feature: Specifying the PKI domain for secure communication with collectors New feature: DRNI MAD New feature: Configuring basic DHCPv6 snooping features on a VXLAN network New feature: Configuring BGP VPNv6 routes to use private network next hops New feature: Configuring BGP route replication between public and VPN instances New feature: Displaying IPv6 multicast group information that MLD snooping learns from EVPN New feature: Enabling password control blacklist for all user line types New feature: Configuring ND |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|---|
| | | | | <p>attack detection for a VSI</p> <ul style="list-style-type: none"> • New feature: MVXLAN <p>Modified feature includes:</p> <ul style="list-style-type: none"> • Modified feature: Configuring LLDP to advertise IPv4 and IPv6 management addresses at the same time • Modified feature: Specifying a padding mode for the Interface-ID option on the DHCPv6 relay agent • Modified feature: Displaying monitor link group information • Modified feature: Displaying the path that the IPv6 packets traverse from source to destination • Modified feature: Exchanging both IPv4 and IPv6 routes with a BGP peer • Modified feature: Enabling load balancing for BGP routes that have different AS_PATH attributes of the same length • Modified feature: Enabling BGP to exchange labeled routes with a peer or peer group • Modified feature: Configuring a static neighbor entry • Modified feature: Modifying information in received BGP EVPN IP prefix advertisement routes on an ED • Modified feature: Configuring passive TCP connection settings for a VXLAN OVSDB VTEP • Modified feature: Setting the password for local password authentication • Modified feature: Enabling an ED to replace the L3 VXLAN ID, RD, and route targets of BGP EVPN routes • Modified feature: Displaying Ethernet service instance information • Modified feature: Configuring OSPF FRR • Modified feature: Displaying routing table information • Modified feature: Configuring the maximum number of Add-Path optimal routes that |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|---|
| | | | | <ul style="list-style-type: none"> can be advertised to all peers Modified feature: Configuring the BGP Additional Paths capabilities Modified feature: Configuring the maximum number of Add-Path optimal routes that can be advertised to a peer or peer group Modified feature: Password control restriction features Modified feature: Flow-mirroring traffic to interfaces Modified feature: Configuring ARP scanning Modified feature: NETCONF logging Modified feature: Enabling route router MAC replacement for a peer or peer group on an EVPN-DCI ED Modified feature: Value range change for the extended community attribute list number in an RR reflection policy Modified feature: Allowing a DR interface to be up when it is the only member of a DR group Modified feature: Configuring a static route Modified feature: Configuring an IPv6 static route Modified feature: Displaying OSPFv3 LSDB information Modified feature: Configuring BGP extended community lists Modified feature: Replicating VLINK direct routes Modified feature: Specifying a controller for an OpenFlow switch and configuring the main connection to the controller Modified feature: Displaying the CPU usage statistics Modified feature: Displaying memory usage information Modified feature: Password complexity checking Modified feature: Configuring kernel thread deadloop detection |

| Version number | Last version | Release Date | Release type | Remarks |
|--------------------|--------------------|--------------|-----------------|---|
| | | | | <ul style="list-style-type: none"> Modified feature: Configuring a traffic redirecting action in a QoS traffic behavior |
| | | | | <p>Added feature includes:</p> <ul style="list-style-type: none"> New features: Fundamentals features New features: Layer 2-LAN switching features New features: Layer 3-IP services features New features: Layer 3-IP routing features New features: IP multicast features New features: MPLS features New features: ACL and QoS features New features: Security features New features: High availability features New features: OpenFlow features New features: VXLAN features New features: EVPN features <p>Modified feature includes:</p> <ul style="list-style-type: none"> Modified feature: Software patching by using issu commands Modified feature: Automatic configuration Modified feature: Collision handling in BFD MAD Modified feature: Processing after the link mode of an Ethernet interface is switched Modified feature: Link state change suppression on an interface Modified feature: Setting the global aging timer for ND entries in stale state Modified feature: Setting the interface-specific aging timer for ND entries in stale state Modified feature: Configuring OSPF area authentication Modified feature: Configuring OSPF interface authentication Modified feature: Configuring a virtual link Modified feature: Displaying statistics for the IGMP messages and PIMv2 hello |
| 5940-CMW710-R 2702 | 5940-CMW710-F27 01 | 2019-06-06 | Release version | |

| Version number | Last version | Release Date | Release type | Remarks |
|-------------------|----------------------|--------------|-----------------|--|
| | | | | <p>messages learned through IGMP snooping</p> <ul style="list-style-type: none"> Modified feature: Displaying statistics for the MLD messages and IPv6 PIM hello messages learned through MLD snooping Modified feature: Enabling a VPN instance to replicate routes from the public network or other VPN instances Modified feature: Using BFD to verify LSP or MPLS TE tunnel connectivity Modified feature: MAC authentication VLAN mode Modified feature: Port security MAC move Modified feature: Creating a BFD session for detecting the local interface state Modified feature: Associating Track with application modules Modified feature: Configuring an interface event for a CLI-defined EAA monitor policy Modified feature: Configuring an EAA monitor policy by using Tcl Modified feature: Removing a TCP or UDP listening service for a VPN instance Modified feature: Specifying the role of the device in the VCF fabric Modified feature: Frame match criteria of VXLAN Ethernet service instances <p>Delete feature includes:</p> <ul style="list-style-type: none"> Deleted feature: Logging NETCONF row operations <p>Fix bugs.</p> |
| 5940-CMW710-F2701 | 5940-CMW710-R2612P02 | 2018-10-31 | Feature version | <p>Added feature includes:</p> <ul style="list-style-type: none"> New feature: Filtering the output from a display command by using multiple conditions New feature: Executing a batch file New feature: Disabling automatic system-wide next-startup configuration file operations New feature: Using an |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|--|
| | | | | <p>encrypted configuration file to roll back configuration</p> <ul style="list-style-type: none"> • New feature: Support for encrypted configuration files for configuration comparison • New feature: ND MAD on management Ethernet ports • New feature: Fast retrain • New feature: Displaying the status and packet statistics of interfaces • New feature: Displaying operating status and information of all interfaces except subinterfaces • New feature: Enabling PFC on all interfaces • New feature: Configuring automatic link aggregation • New feature: Disabling the default action of selecting a Selected port for dynamic aggregation groups that have not received LACPDUs • New feature: Physical state change suppression on aggregate interfaces • New feature: VXLAN multiport ARP entries • New feature: Enabling recording user port move for ARP • New feature: Enabling recording user IP address conflicts • New feature: Enabling interface consistency check for ARP and MAC address entries • New feature: Specifying a DHCP relay address pool for DHCP clients • New feature: Configuring the padding mode for the Vendor-Specific sub-option • New feature: Configuring IPv6 ND direct route advertisement • New feature: Configuring ND snooping in a VXLAN • New feature: Enabling recording user IPv6 address conflicts • New feature: Enabling recording user port move for ND • New feature: Enabling ND |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|--|
| | | | | <p>logging for user online and offline events</p> <ul style="list-style-type: none"> • New feature: Specifying the boot file URL in RA messages • New feature: Setting the aging timer for ND entries in stale state on an interface • New feature: Specifying DNS server information in RA messages • New feature: Specifying the DNS domain name suffix information in RA messages • New feature: Suppressing DNS information in RA messages • New feature: Enabling the DHCPv6 relay agent to advertise host routes for assigned IPv6 addresses • New feature: Enabling IPv6 address binding conversion for IP source guard • New feature: Configuring DHCPv6 guard • New feature: Ignoring router IDs during optimal route selection • New feature: Redistributing the local network specified in the public instance or a VPN instance • New feature: Apply route update interval setting to withdrawn routes • New feature: Enabling prioritized withdrawal of specific routes • New feature: Configuring BGP RPKI • New feature: Configuring BGP policy accounting • New feature: Configuring BGP to immediately send updates for routes that match an IPv6 prefix list • New feature: BGP optimal route selection delay • New feature: Specifying a VPN instance for the BMP server • New feature: Routing policy change delay • New feature: Support of a routing policy for L3VNI configuration |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|---|
| | | | | <ul style="list-style-type: none"> • New feature: Configuring the BGP RPKI validation state match criterion for a routing policy node • New feature: Enabling logging for source MAC-based ARP attack detection • New feature: Enabling ND attack detection logging • New feature: Configuring BFD session flapping suppression • New feature: Configuring the timer that delays reporting the first BFD session establishment failure to the data link layer • New feature: Enabling special processing for BFD sessions • New feature: SBFD • New feature: Associating Track with a control-mode BFD session • New feature: Enabling the device to send IGMP general queries upon a path change • New feature: Enabling the device to send MLD general queries upon a path change • New feature: Enabling embedded RP • New feature: Specifying a local network to be advertised in the public instance or a VPN instance • New feature: Changing the next hop address of VPNv4 routes to a VPN address • New feature: Displaying tunnel policy information • New feature: SBFD for MPLS • New feature: MPLS SR • New feature: Displaying buffer usage statistics for interfaces • New feature: Configuring an EAP profile • New feature: AAA test • New feature: Displaying and clearing history authentication and accounting load statistics for all RADIUS servers • New feature: 802.1X offline detection |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|---|
| | | | | <ul style="list-style-type: none"> • New feature: 802.1X unauthenticated user aging • New feature: 802.1X EAP-TLS fragmentation for packets sent to the server • New feature: User aging for unauthenticated MAC authentication users • New feature: MAC authentication method • New feature: User-specific MAC authentication offline detection • New feature: NAS-Port-Type attribute • New feature: Checking the issuing of category-2 portal filtering rules • New feature: Configuring user authentication timeout timer for global password control • New feature: Configuring password control for local network access users • New feature: Configuring password change at first login • New feature: Enabling the SSL server to send the complete certificate chain during SSL negotiation • New feature: Support of RRPP for the link-up delay timer • New feature: gRPC • New feature: Protocol buffers • New feature: Specifying the NTP time-offset thresholds for log and trap outputs • New feature: Specifying the SNTP time-offset thresholds for log and trap outputs • New feature: Ignoring the permit flag added by OpenFlow • New feature: ND flood suppression • New feature: IP address configuration of the VTEPs in a DR system • New feature: Redistribution of MAC/IP advertisement routes into BGP unicast routing tables • New feature: Route exchange between the BGP |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|---|
| | | | | <p>EVPN address family and the BGP VPNv4 or VPNv6 address family</p> <ul style="list-style-type: none"> • New feature: Disabling the ACLs issued by the OVSDB controller • New feature: Replacement of the L3 VXLAN ID and RD of IP prefix advertisement routes on EDs • New feature: BGP EVPN route advertisement suppression • New feature: Removal of the default-gateway extended community attribute from EVPN gateway routes • New feature: Configuring a 16/32-Gbps FC interface • New feature: Enabling area ID-to-F_Port binding • New feature: BB_Credit recovery for FC interfaces <p>Modified feature includes:</p> <ul style="list-style-type: none"> • Modified feature: Software upgrade • Modified feature: Using issu commands for an ISSU • Modified feature: Per-packet load sharing mode of aggregation groups • Modified feature: Configuring MAC-to-VLAN entries • Modified feature: Configuring the advertisable TLVs • Modified feature: Specifying the management address advertised in global management address TLV advertisement setting • Modified feature: Configuring aging timer for dynamic ARP entries • Modified feature: Enabling ARP snooping • Modified feature: Displaying ARP snooping entries • Modified feature: Clearing ARP snooping entries • Modified feature: Configuring DHCP snooping handling strategy for Option 82 in request messages • Modified feature: Setting the interface MTU for IPv6 packets • Modified feature: Support for |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|---|
| | | | | <ul style="list-style-type: none"> specifying next hop indexes in static routes Modified feature: Support for configuring OSPF commands in VSI interface view Modified feature: Displaying IS-IS LSP log information Modified feature: Clearing IS-IS LSP log information Modified feature: Filtering advertised/received BGP routes Modified feature: Filtering routes advertised to or received from a peer or peer group by using an ACL Modified feature: Configuring BMP Modified feature: Displaying detailed BGP routing information Modified feature: Applying a routing policy to routes outgoing to a peer or peer group Modified feature: Specifying an ACL Modified feature: Defining an ACL match criterion in a traffic class of a QoS policy Modified feature: Applying a QoS policy globally Modified feature: Configuring a test profile for RADIUS server status detection Modified feature: RADIUS server quiet timer Modified feature: Specifying the source IP address for outgoing RADIUS packets Modified feature: Specifying the source IP address for outgoing HWTACACS packets Modified feature: Including user IP addresses in MAC authentication requests Modified feature: Configuring MAC-based MAC authentication user accounts Modified feature: Handling new MAC access attempts in a VLAN after port security's MAC address limit for that VLAN is reached Modified feature: Port |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|--|
| | | | | <p>security NTK feature</p> <ul style="list-style-type: none"> • Modified feature: Password handling manners with password control enabled globally • Modified feature: RSA key modulus length • Modified feature: Key modulus length of the RSA key pair used for certificate request in a PKI domain • Modified feature: Configuring the ECDSA signature authentication method in an IKE proposal • Modified feature: Specifying the DH group used for key negotiation in IKE phase 1 • Modified feature: Keyboard-interactive authentication support for SSH users • Modified feature: Displaying IPv4 source guard bindings • Modified feature: Displaying IPv6 source guard bindings • Modified feature: Configuring ARP attack detection logging • Modified feature: Creating a BFD session for detecting the local interface state • Modified feature: Configuring the BFD authentication mode for single-hop BFD control packets • Modified feature: Setting the minimum interval for receiving BFD echo packets • Modified feature: Support for specific BFD commands in VSI interface view • Modified feature: Setting the maximum number of active routes in a VPN instance • Modified feature: Enabling BFD on an MPLS TE tunnel interface • Modified feature: Setting the length of ICMP or ICMPv6 echo requests for the ping operation • Modified feature: Specifying a source interface for NTP messages • Modified feature: PTP • Modified feature: Shutting down an interface by using |

| Version number | Last version | Release Date | Release type | Remarks |
|-----------------------|----------------------|--------------|-----------------|--|
| | | | | <p>OpenFlow</p> <ul style="list-style-type: none"> Modified feature: Configuring flow sampling of sFlow Modified feature: Configuring counter sampling of sFlow Modified feature: Testing the reachability of remote VMs in VXLANs Modified feature: Enabling packet statistics for automatically created VXLAN tunnels Modified feature: Disabling flooding for a VSI Modified feature: Setting the maximum bandwidth for an AC Modified feature: Displaying information about peers automatically discovered through BGP in EVPN Modified feature: PW redundancy Modified feature: Displaying log buffer information and buffered logs Modified feature: Configuring the timestamp for logs sent to log hosts Modified feature: Configuring log output to a log host Modified feature: IRF master election during automated VCF fabric deployment <p>Fix bugs.</p> |
| 5940-CMW710-R-2612P02 | 5940-CMW710-R2612P01 | 2018-09-27 | Release version | <p>Added feature includes:</p> <ul style="list-style-type: none"> New feature: Default VXLAN decapsulation <p>Modified features includes:</p> <ul style="list-style-type: none"> Modified feature: Setting the maximum bandwidth for a VSI Modified feature: Setting the broadcast, multicast, or unknown unicast restraint bandwidth for a VSI Modified feature: Configuring a traffic policing action in a traffic behavior <p>Fix bugs.</p> |
| 5940-CMW710-R-2612P01 | 5940-CMW710-R2612 | 2018-08-24 | Release version | <p>Added feature includes:</p> <ul style="list-style-type: none"> New feature: Configuring the DHCP relay agent to forward DHCP replies based on MAC address table <p>Fix bugs.</p> |

| Version number | Last version | Release Date | Release type | Remarks |
|--------------------|-------------------|--------------|-----------------|---|
| 5940-CMW710-R-2612 | 5940-CMW710-R2610 | 2018-06-01 | Release version | <p>Added feature includes:</p> <ul style="list-style-type: none"> • New feature: Load sharing mode for tunneled traffic on aggregate links • New feature: Interfaces excluded from the shutdown action by DRNI MAD • New feature: Enabling split horizon forwarding • New feature: Enabling remote fault signal detection • New feature: Setting PFC thresholds • New feature: Packet statistics on Layer 3 Ethernet subinterfaces • New feature: Packet statistics on Layer 3 aggregate subinterfaces • New feature: Simple multichassis link aggregation (S-MLAG) • New feature: DHCPv6 relay agent support for Option 79 • New feature: Setting the action that drops matching packets when all next hops specified on an IPv4 or IPv6 policy node are invalid • New feature: Setting the SoO extended community attribute for BGP routes • New feature: Displaying multicast group information that IGMP snooping learns from EVPN • New feature: Layer 2 multicast forwarding mode • New feature: Layer 2 IPv6 multicast forwarding mode • New feature: Setting the maximum number of cached unknown multicast packets • New feature: Setting the maximum number of cached unknown IPv6 multicast packets • New feature: Disabling generation of IP prefix advertisement routes for the subnets of a VSI interface • New feature: EVPN multihoming • New feature: EVPN multicast • New feature: Packet statistics for manually created VXLAN and |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|--|
| | | | | <p>VXLAN-DCI tunnels</p> <ul style="list-style-type: none"> • New feature: Packet statistics for automatically created VXLAN tunnels • New feature: EVPN-DCI dual-homing • New feature: PTP • New feature: Support of OpenFlow for matching VSI interfaces and tunnel interfaces <p>Modified features includes:</p> <ul style="list-style-type: none"> • Modified feature: Resource monitoring • Modified feature: Displaying information about dropped packets on interfaces • Modified feature: Configuring port isolation • Modified feature: Hash seed for the global link aggregation load sharing mode • Modified feature: Displaying the configuration and running status of loop detection • Modified feature: Displaying LLDP neighbor information • Modified feature: Displaying DRNI configuration consistency check information • Modified feature: DR system auto-recovery in DRNI • Modified feature: Displaying ARP entries • Modified feature: Setting the DHCP server response timeout time for DHCP server switchover • Modified feature: Displaying IPv4/IPv6 route statistics • Modified feature: Clearing route statistics • Modified feature: Route-type match criterion • Modified feature: Route redistribution for OSPF • Modified feature: Route redistribution for IS-IS • Modified feature: Creating a summary route in the BGP routing table • Modified feature: Specifying an AS path list to filter routes incoming from or outgoing to a peer or peer group |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|--|
| | | | | <ul style="list-style-type: none"> Modified feature: Specifying an ACL to filter routes incoming from or outgoing to a peer or peer group Modified feature: Specifying a prefix list to filter routes incoming from or outgoing to a peer or peer group Modified feature: Specifying a routing policy to filter routes incoming from or outgoing to a peer or peer group Modified feature: Interval for sending the same update to a BGP peer or peer group Modified feature: Route redistribution for OSPFv3 Modified feature: IGMP snooping proxying Modified feature: Displaying information about dynamic IGMP snooping group entries Modified feature: Displaying information about Layer 2 IP multicast groups Modified feature: MLD snooping proxying Modified feature: Displaying information about dynamic MLD snooping group entries Modified feature: Displaying information about Layer 2 IPv6 multicast groups Modified feature: Setting the maximum size of a PIM join or prune message Modified feature: Setting the maximum size of an IPv6 PIM join or prune message Modified feature: Displaying ACL configuration and match statistics for Layer 2 ACLs Modified feature: Displaying ACL configuration and match statistics for user-defined ACLs Modified feature: QoS policy Modified feature: Configuring Layer 3 remote port mirroring Modified feature: ARP scanning Modified feature: Configuring a tunnel policy Modified feature: Specifying a BFD session parameter template for the MPLS TE backup tunnel |

| Version number | Last version | Release Date | Release type | Remarks |
|--------------------|----------------------|--------------|-----------------|--|
| | | | | <ul style="list-style-type: none"> Modified feature: Setting the minimum interval for receiving BFD echo packets Modified feature: Shutting down an interface by OpenFlow Modified feature: Enabling packet loss prevention for OpenFlow forwarding Modified feature: Configuring flow tables for an OpenFlow instance Modified feature: Automated underlay network deployment of VCF fabric Deleted features includes: <ul style="list-style-type: none"> Removed feature: EEE on an Ethernet interface Removed feature: Outgoing packets filtering on a portal-enabled interface Fix bugs. |
| 5940-CMW710-R-2610 | 5940-CMW710-R2609 | 2017-11-30 | Release version | Added feature: <ul style="list-style-type: none"> New feature: MACsec support on the LSWM18CQMSEC(JH957A) interface card New feature: Cipher suite for MACsec encryption Modified features includes: <ul style="list-style-type: none"> Modified feature: Displaying recommended ISSU methods Modified feature: Specifying a VPN instance for DRNI keepalive packets Modified feature: Creating ACLs and entering ACL view Modified feature: Displaying QoS and ACL resource usage Modified feature: Displaying the configuration and statistics for traffic policing configured by using a QoS policy Modified feature: Setting the maximum link bandwidth and maximum reservable bandwidth for MPLS TE traffic Fix bugs. |
| 5940-CMW710-R-2609 | 5940-CMW710-R2608P02 | 2017-11-24 | Release version | Modified features includes: <ul style="list-style-type: none"> Modified feature: QPPB Fix bugs. |

| Version number | Last version | Release Date | Release type | Remarks |
|-----------------------|-------------------|--------------|-----------------|--|
| 5940-CMW710-R-2608P02 | 5940-CMW710-F2608 | 2017-10-30 | Release version | <p>Added feature:</p> <ul style="list-style-type: none"> • New feature: Automatic obtaining of the login username for temporary user role authorization • New feature: Binary-only configuration saving mode • New feature: Remote configuration archiving • New feature: Automatic member port assignment of dynamic aggregation groups • New feature: Port speed ignorance in setting the aggregation states of member ports • New feature: DRNI configuration consistency check • New feature: Enabling route fast switchover • New feature: Disabling BGP session establishment • New feature: DSCP value for outgoing IGMP snooping protocol packets • New feature: DSCP value for outgoing MLD snooping protocol packets • New feature: DSCP value for outgoing IGMP protocol packets • New feature: DSCP value for outgoing MLD protocol packets • New feature: DSCP value for outgoing IPv6 PIM protocol packets • New feature: DSCP value for outgoing MSDP protocol packets • New feature: DSCP value for outgoing multicast VPN protocol packets • New feature: SSH algorithm renegotiation and key re-exchange • New feature: Source IP address for outgoing SCP packets • New feature: Setting the port shutdown mode for an interface • New feature: Configuring prefix SID-based static SR • New feature: Adding an outer VLAN tag for incoming |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|--|
| | | | | <p>packets</p> <ul style="list-style-type: none"> • New feature: Applying an ACL to a list of VLAN interfaces for packet filtering • New feature: Clearing queue-based outgoing traffic statistics for interfaces • New feature: QPPB • New feature: Forcibly sending RADIUS stop-accounting packets • New feature: The escape critical VSI feature of port security • New feature: Connection recording policy • New feature: Displaying versions of the template file for automated VCF fabric provisioning • New feature: VLAN tag processing rules • New feature: Distributed VXLAN IPv6 gateways • New feature: Local-MAC address learning • New feature: EVPN support for IPv6 sites • New feature: EVPN distributed relay support for VXLAN tunnel IPL • New feature: EVPN global MAC address <p>Modified features includes:</p> <ul style="list-style-type: none"> • Modified feature: FTP server login • Modified feature: Hardware resource mode • Modified feature: Free-memory monitoring • Modified feature: Removing support for displaying and clearing information about packets dropped by subinterfaces • Modified feature: DRNI • Modified feature: Setting the dynamic ARP learning limit for an interface • Modified feature: Setting the aging time for dynamic ARP entries • Modified feature: Displaying information about DHCP address pools • Modified feature: Specifying |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|--|
| | | | | <p>DHCP servers on the DHCP relay agent</p> <ul style="list-style-type: none"> • Modified feature: Displaying DHCP server addresses • Modified feature: Associating the output interface of a static route with a track entry • Modified feature: Setting the time to wait for the End-of-RIB marker • Modified feature: Multicast VPN • Modified feature: Configuring seed value and shift value in load sharing algorithm • Modified feature: Advertising OSPF link state information to BGP • Modified feature: Creating a BFD session for detecting the local interface state • Modified feature: Configuring an IPv4 advanced ACL rule • Modified feature: Configuring a Layer 2 ACL rule • Modified feature: Defining an ACL match criterion • Modified feature: Configuring an accounting action • Modified feature: Associating a traffic class with a traffic behavior in a QoS policy • Modified feature: Configuring WFQ queuing parameters for an interface • Modified feature: Configuring queuing parameters in a queue scheduling profile • Modified feature: Applying the QoS policy to an interface • Modified feature: MAC authentication critical VSI • Modified feature: Displaying IPv4SG bindings • Modified feature: Displaying statistics for ARP attack detection • Modified feature: Clearing statistics for ARP attack detection • Modified feature: ARP attack detection logging • Modified feature: Setting the broadcast, multicast, or unknown unicast restraint bandwidth for a VSI • Modified feature: EVPN |

| Version number | Last version | Release Date | Release type | Remarks |
|-------------------|-------------------|--------------|-----------------|--|
| | | | | distributed relay |
| | | | | <ul style="list-style-type: none"> Modified feature: Configuring the Ethernet OAM connection detection timers |
| | | | | Fix bugs. |
| | | | | Added feature: |
| | | | | <ul style="list-style-type: none"> New feature: Management Ethernet interface backup New feature: S-channel bundle interface New feature: Dispute guard New feature: Displaying the number of OpenFlow ARP entries New feature: Specifying the source IPv6 address for relayed DHCPv6 requests New feature: Enabling recording IPv6 prefix information in DHCPv6 snooping entries New feature: Enabling OSPF to limit the LSU transmit rate New feature: Configuring BGP to first send updates of the default route New feature: Setting the BGP route sending rate New feature: Setting the session retry timer New feature: Static SR over MPLS New feature: 802.1X user logging New feature: MAC address information display for 802.1X users in 802.1X VLANs of a specific type New feature: MAC authentication support for blackhole MAC attribute assignment New feature: MAC authentication user logging New feature: MAC address information display for MAC authentication users in MAC authentication VLANs of a specific type New feature: Displaying IPv6SG prefix bindings New feature: Specifying the value type for the IfIndex or vrfindex element New feature: Resource monitoring |
| 5940-CMW710-F2608 | 5940-CMW710-F2606 | 2017-07-12 | Feature version | |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|---|
| | | | | <ul style="list-style-type: none"> • New feature: Configuring routing policy-based recursive lookup for IPv6 BGP routes • New feature: Setting the maximum number of active routes • New feature: BGP additional path • New feature: IGMP snooping querier election • New feature: MLD snooping querier election • New feature: Stateful ACL • New feature: Fabric Monitor and Analyzer • New feature: Support for DSCP values in SNMP responses • New feature: Setting the DSCP value in information center log packets sent to a log host • New feature: Software-based MAC learning on an interface • New feature: VXLAN over VXLAN • New feature: Packet statistic collection for Ethernet service instances of a VLAN • New feature: 802.1X VSI manipulation • New feature: MAC authentication VSI manipulation • New feature: Dynamic Ethernet service instance mappings of VSIs • New feature: Specifying the Telnet service port number <p>Modified features includes:</p> <ul style="list-style-type: none"> • Modified feature: ARP attack detection • Modified feature: NETCONF logging • Modified feature: Performing <edit-config> operations during the NETCONF configuration rollback process • Modified feature: NETCONF CLI operations • Modified feature: ARP MAD collision handling process • Modified feature: Setting the device operating mode to |

| Version number | Last version | Release Date | Release type | Remarks |
|-------------------|-------------------|--------------|-----------------|---|
| | | | | PEX <ul style="list-style-type: none"> Modified feature: Applying an ACL to an interface for packet filtering Modified feature: Displaying ACL application information for packet filtering Modified feature: Applying a QoS policy to an interface Modified feature: Displaying QoS policies applied to interfaces Modified feature: NTP polling intervals Modified feature: Maximum number of port isolation groups Modified feature: Specifying startup image files Modified feature: Monitoring CPU usage Modified feature: Aggregation member port support for service instances bound to a cross connect or VSI Modified feature: Displaying routing table information Modified feature: Displaying route statistics Modified feature: Displaying BGP peer or peer group information Modified feature: Configuring label allocation Modified feature: 802.1X Modified feature: MAC authentication Modified feature: Port security Modified feature: TLS version used for the OVSDB server to establish an SSL connection to the controller |
| 5940-CMW710-F2606 | 5940-CMW710-F2605 | 2017-04-28 | Feature version | Added feature: <ul style="list-style-type: none"> Transparent LACPDU transmission Configuring a description for a QoS traffic class |
| 5940-CMW710-F2605 | 5940-CMW710-F2604 | 2017-04-07 | Feature version | Added feature: <ul style="list-style-type: none"> Traffic statistics collection for a Layer 3 Ethernet subinterface Exclusion of a subnet from load sharing on link aggregations |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------|--------------|--------------|--------------|---|
| | | | | <ul style="list-style-type: none"> Releasing SSH connections Gratuitous ARP packet retransmission for the device MAC address change Configuring DHCP snooping for VLANs Configuring basic DHCP snooping features in a VXLAN network HTTP redirect Enabling the compressed ECMP mode Periodic sending of ARP requests to the next hops of static routes Disabling route recursion policy control for routes received from a peer or peer group Applying an ACL to an Ethernet service instance for packet filtering Displaying diagnostic information about applied QoS policies Configuring the switch to trust the 802.1p priority in the outer IP header of VXLAN packets AAA support for VSI-based ACL resource allocation Quiet timer for users logged off by the authorization-fail-offline feature Port security's limit on the number of MAC addresses for specific VLANs on a port Displaying the role of the device in the VCF fabric Configuring NETCONF to use module-specific namespaces Support of OpenFlow for specifying controllers to which ARP packets are forbidden to be reported Reachability test for remote user terminals in VXLANs VLAN-based VXLAN assignment <p>Modified features includes:</p> <ul style="list-style-type: none"> Enabling link flapping protection on an interface Setting the aging time for dynamic ARP entries |

| Version number | Last version | Release Date | Release type | Remarks |
|-------------------|-------------------|--------------|-----------------|---|
| | | | | <ul style="list-style-type: none"> • Displaying ARP entries • Enabling a BGP instance • Displaying information about dynamic IGMP snooping group entries • Displaying information about dynamic MLD snooping group entries • Displaying VSI information • Creating an ACL • Displaying ACL application information for packet filtering • Displaying ACL packet filtering statistics • Displaying ACL application details for packet filtering • Clearing the packet filtering statistics (including the accumulated statistics) for an ACL • Creating a rule for a user-defined ACL • MAC authentication support for ACL assignment • Specifying a RADIUS or HWTACACS server • Maximum number of VSANs supported • Setting free-memory thresholds • Setting the maximum size for the information center trace log file • VXLAN tunnel • Interfaces for connecting EDs of an EVPN-DCI network <p>Deleted features:</p> <ul style="list-style-type: none"> • Configuring the switch to trust the DSCP priority in the outer IP header of VXLAN packets <p>Fixed bugs.</p> |
| 5940-CMW710-F2604 | 5940-CMW710-F2603 | 2017-02-28 | Feature version | <p>Added feature:</p> <ul style="list-style-type: none"> • New features: Fundamentals features • New features: Layer 2-LAN switching features • New features: Layer 3-IP services features • New features: Layer 3-IP routing features • New features: IP multicast features • New features: MPLS features |

| Version number | Last version | Release Date | Release type | Remarks |
|----------------------|----------------------|--------------|-----------------|---|
| | | | | <ul style="list-style-type: none"> • New features: ACL and QoS features • New features: Security features • New features: High availability features • New features: VXLAN features • New features: EVPN features Modified features: <ul style="list-style-type: none"> • Modified feature: Username format modification for device login • Modified feature: Interval for an OpenFlow instance to reconnect to a controller • Modified feature: Displaying group table statistics • Modified feature: Setting the voice VLAN aging timer • Modified feature: Creating a VLAN Fixed bugs. |
| 5940-CMW710-F2603 | 5940-CMW710-R2509P02 | 2016-12-19 | Feature version | None |
| 5940-CMW710-R2509P02 | 5940-CMW710-R2509P01 | 2016-11-03 | Release version | None |
| 5940-CMW710-R2509P01 | 5940-CMW710-R2509 | 2016-09-18 | Release version | None |
| 5940-CMW710-R2509 | 5940-CMW710-R2508 | 2016-09-07 | Release version | None |
| 5940-CMW710-R2508 | 5940-CMW710-R2507 | 2016-07-29 | Release version | None |
| 5940-CMW710-R2507 | 5940-CMW710-R2506 | 2016-07-14 | Release version | None |
| 5940-CMW710-R2506 | First release | 2016-05-30 | Release version | None |

Hardware and software compatibility matrix



CAUTION:

To avoid an upgrade failure, use [Table 2](#) to verify the hardware and software compatibility before performing an upgrade.

Table 2 Hardware and software compatibility matrix

| Item | Specifications |
|----------------|-----------------|
| Product family | HPE 5940 Series |

| Item | Specifications |
|---|---|
| Hardware platform | FF 5940 32QSFP+ JH396A FF 5940 48XGT 6QSFP28 JH391A FF 5940 48SFP+ 6QSFP28 JH390A FF 5940 2-slot Switch JH397A FF 5940 4-slot Switch JH398A |
| Memory | 4GB |
| Flash | 1GB |
| Boot ROM version | Version 241 or higher (Note: Perform the command display version command in any view to view the version information. Please see Note ②) |
| Software images and their MD5 checksums | 5940-CMW710-R6710P03.ipe: 726dcdea68fdb5bb83a4b042eeac23a9 |
| iMC version | iMC PLAT 7.3 (E0706P13) iMC EAD 7.3 (E0623) iMC MVM 7.3 (E0511) iMC QoSM 7.3 (E0506) iMC VXLAN 7.3 (E0706P13) iMC BIMS 7.3 (E0509H01) UCenter SHM 7.3 (E0711) UCenter NTA 7.3 (E0711) iMC EIA 7.3 (E0627) |
| iNode version | iNode(E0595) |
| VCF Controller version | E2187 E2180P11 |
| Web version | None |
| OAA version | None |

To display version information for the system software and Boot ROM of 5940:

```
<HPE>display version
```

```
HPE Comware Software, Version 7.1.070, Release 6710P03          ----- Note①
```

```
Copyright (c) 2010-2023 Hewlett-Packard Development Company, L.P.
```

```
HPE FF 5940-32QSFP+ Switch uptime is 0 weeks, 0 days, 18 hours, 23 minutes
```

```
Last reboot reason : USER reboot
```

```
Boot image: flash:/5940-cmw710-boot-r6710p03.bin
```

```
Boot image Version: 7.1.070, Release 6710P03
```

```
Compiled May 05 2023 11:00:00
```

```
System image: flash:/5940-cmw710-system-r6710p03.bin
```

```
System image Version: 7.1.070, Release 6710P03
```

```
Compiled May 05 2023 11:00:00
```

```
Slot 1:
```

```
Uptime is 0 weeks,0 days,18 hours,23 minutes
```

```

FF 5940-32QSFP+ Switch with 1 RMI XLP208 Processors
BOARD TYPE:          FF 5940 32QSFP+ Switch
DRAM:                4096M bytes
FLASH:               1024M bytes
PCB 1 Version:       VER.A
PCB 2 Version:       VER.A
Bootrom Version:     241
CPLD 1 Version:      002
CPLD 2 Version:      004
Release Version:     HPE FF 5940-32QSFP+ Switch-6710P03
Patch Version  :     None
Reboot Cause  :      UserReboot
[SubSlot 0] 32QSFP Plus
----- Note②

```

ISSU compatibility list

ISSU provides two upgrade types: compatible upgrade and incompatible upgrade. [Table 3](#) provides the approved ISSU upgrade types only between the current version and the history versions within the past 18 months. This matrix does not include history versions that are 18 months earlier than the current version, for which, no ISSU upgrade verification is performed.

For more information about ISSU, see the fundamental configuration guide for the device.

Table 3 ISSU compatibility list

| Current version | History version | ISSU upgrade method |
|----------------------|----------------------|---------------------|
| 5940-CMW710-R6710P03 | 5940-CMW710-R6710P01 | Compatible |
| | 5940-CMW710-R6710 | Compatible |
| | 5940-CMW710-F6705 | Compatible |
| | 5940-CMW710-E6702 | Compatible |
| | 5940-CMW710-R2717 | Uncompatible |
| | 5940-CMW710-R2702 | Uncompatible |
| | 5940-CMW710-F2701 | Uncompatible |
| | 5940-CMW710-R2612P02 | Uncompatible |
| | 5940-CMW710-R2612P01 | Uncompatible |
| | 5940-CMW710-R2612 | Uncompatible |
| | 5940-CMW710-R2610 | Uncompatible |
| | 5940-CMW710-R2609 | Uncompatible |
| | 5940-CMW710-R2608P02 | Uncompatible |
| | 5940-CMW710-F2608 | Uncompatible |
| | 5940-CMW710-F2606 | Uncompatible |
| | 5940-CMW710-F2605 | Uncompatible |
| | 5940-CMW710-F2604 | Uncompatible |
| | 5940-CMW710-F2603 | Uncompatible |
| | 5940-CMW710-R2510P02 | Uncompatible |

| | |
|----------------------|--------------|
| 5940-CMW710-R2510P01 | Uncompatible |
| 5940-CMW710-R2509P02 | Uncompatible |
| 5940-CMW710-R2509P01 | Uncompatible |
| 5940-CMW710-R2509 | Uncompatible |
| 5940-CMW710-R2508 | Uncompatible |
| 5940-CMW710-R2507 | Uncompatible |
| 5940-CMW710-R2506 | Uncompatible |

Upgrading restrictions and guidelines

- 1、 Before performing a software upgrade, it is important to refer to the *Software Feature Changes* document for any feature changes in the new version. Also check the most recent version of the related documentation (see "[Related documentation](#).") available for more information about feature configuration and commands.
- 2、 Among the 5940 Switch Series, the version F2701 and later might fail to operate properly on switches with the product codes 5940 48XGT 6QSFP+ JH394A and 5940 48SFP+ 6QSFP+ JH395A due to memory limitations. For these models, as a best practice, use the version R2612P02.

Hardware feature updates

R6710P03

None.

R6710P01

None.

R6710

None.

F6705

None.

E6702

None.

R2717

None.

R2702

None.

F2701

Supports the following switches and modules:

- HPE 5940 16-port SFP28 Converged Module JH339A

R2612P02

None.

R2612P01

None.

R2612

None.

R2610

None.

R2609

None.

R2608P02

None.

F2608

None.

F2606

None.

F2605

None.

F2604

None.

F2603

None.

R2509P02

Supports the following switches and modules:

- HPE 5940 2-port QSFP+ and 2-port QSFP28 Module JH409A

R2509P01

None.

R2509

None.

R2508

Supports the following switches and modules:

- HPE FlexFabric 5940 2-slot Switch JH397A
- HPE FlexFabric 5940 4-slot Switch JH398A
- HPE 5930 8-port QSFP+ Module JH183A
- HPE 5930 24-port 10GBASE-T and 2-port QSFP+ with MACsec Module JH182A
- HPE 5930 24-port SFP+ and 2-port QSFP+ with MACsec Module JH181A
- HPE 5930 24-port SFP+ and 2-port QSFP+ Module JH180A

R2507

None.

R2506

First release.

Software feature and command updates

For more information about the software feature and command update history, see HPE 5940-CMW710-R6710P03 Release Notes (Software Feature Changes).

MIB updates

Table 4 MIB updates

| Item | MIB file | Module | Description |
|-----------------------------|--|--|---|
| 5940-CMW710-R6710P03 | | | |
| New | None | None | None |
| Modified | None | None | None |
| 5940-CMW710-R6710P01 | | | |
| New | First release | First release | First release |
| Modified | First release | First release | First release |
| 5940-CMW710-R6710 | | | |
| New | None | None | None |
| Modified | None | None | None |
| 5940-CMW710-F6705 | | | |
| New | None | None | None |
| Modified | None | None | None |
| 5940-CMW710-E6702 | | | |
| New | None | None | None |
| Modified | None | None | None |
| 5940-CMW710-R2717 | | | |
| New | None | None | None |
| Modified | None | None | None |
| 5940-CMW710-R2702 | | | |
| New | None | None | None |
| Modified | rfc2674-pbridge.mib | P-BRIDGE-MIB | Modified dot1dPortOutboundAccess PriorityTable |
| 5940-CMW710-F2701 | | | |
| New | None | None | None |
| Modified | None | None | None |
| 5940-CMW710-R2612P02 | | | |
| New | hh3c-l2vpn.mib hh3c-tunnel.mib hh3c-evpn.mib | HH3C-L2VPN-MIB HH3C-TUNNEL-MIB HH3C-EVPN-MIB | Added hh3cL2vpnLinkTable hh3cL2vpnAcCfgTable hh3cL2vpnPwCfgTable Added HH3C-TUNNEL-MIB Added HH3C-EVPN-MIB |
| Modified | None | None | None |

| Item | MIB file | Module | Description |
|-----------------------------|---|--|---|
| 5940-CMW710-R2612P01 | | | |
| New | None | None | None |
| Modified | None | None | None |
| 5940-CMW710-R2612 | | | |
| New | None | None | None |
| Modified | None | None | None |
| 5940-CMW710-R2610 | | | |
| New | None | None | None |
| Modified | None | None | None |
| 5940-CMW710-R2609 | | | |
| New | None | None | None |
| Modified | None | None | None |
| 5940-CMW710-R2608P02 | | | |
| New | hh3c-cbqos2.mib hh3c-8021x-ext2.mib Hh3c-port-security.mib ieee8021x.mib ieee8021x-pae.mib rfc5676-syslog-msg.mib hh3c-mpls-vpn-bgp.mib | HH3C-CBQOS2-MIB HH3C-8021X-EXT2-MIB HH3C-PORT-SECURITY-MIB IEEE8021-PAE-MIB IEEE8021X-PAE-MIB SYSLOG-MSG-MIB HH3C-MPLS-VPN-BGP-MIB | New hh3cCBQoSCarClassName hh3cCBQoSAccountingPktps hh3cCBQoSAccountingPktps New HH3C-8021X-EXT2-MIB HH3C-PORT-SECURITY-MIB IEEE8021-PAE-MIB IEEE8021X-PAE-MIB SYSLOG-MSG-MIB New HH3C-MPLS-VPN-BGP-MIB |
| Modified | rfc5602-pw-mpls-std.mib rfc4273-bgp4.mib | PW-MPLS-STD-MIB BGP4-MIB | Modified pwMplsStorageType Modified bgpPeerConnectRetryInterval bgpPeerHoldTimeConfigured bgpPeerKeepAliveConfigured |
| 5940-CMW710-F2608 | | | |
| New | h3c-lsw-dev-adm.mib hh3c-if-ext.mib | HH3C-LSW-DEV-ADM-MIB HH3C-IF-EXT-MIB | New hh3cLswSubslotSerialNumber hh3cLswSlotSerialNumber New |

| Item | MIB file | Module | Description |
|-----------------------------|----------------------|----------------------|---|
| | | | hh3clfdescription |
| Modified | None | None | None |
| 5940-CMW710-F2606 | | | |
| New | None | None | None |
| Modified | None | None | None |
| 5940-CMW710-F2605 | | | |
| New | None | None | None |
| Modified | None | None | None |
| 5940-CMW710-F2604 | | | |
| New | None | None | None |
| Modified | None | None | None |
| 5940-CMW710-F2603 | | | |
| New | None | None | None |
| Modified | None | None | None |
| 5940-CMW710-R2509P02 | | | |
| New | None | None | None |
| Modified | None | None | None |
| 5940-CMW710-R2509P01 | | | |
| New | None | None | None |
| Modified | None | None | None |
| 5940-CMW710-R2509 | | | |
| New | None | None | None |
| Modified | None | None | None |
| 5940-CMW710-R2508 | | | |
| New | None | None | None |
| Modified | None | None | None |
| 5940-CMW710-R2507 | | | |
| New | None | None | None |
| Modified | lldp-ext-dot1-v2.mib | LLDP-EXT-DOT1-V2-MIB | Modified lldpV2Xdot1ConfigVlanNameTxEnable lldpV2Xdot1ConfigProtoVlanTxEnable |
| 5940-CMW710-R2506 | | | |
| New | First release | First release | First release |
| Modified | First release | First release | First release |

Operation changes

Operation changes in R6710P03

- [202306132304] Added the dropped incoming and outgoing packet counts for aggregate interfaces to the ifmgr/statistics table
Before modification: The ifmgr/statistics table does not provide the dropped incoming and outgoing packet counts for aggregate interfaces.
After modification: The ifmgr/statistics table provides the dropped incoming and outgoing packet counts for aggregate interfaces.
- [202302250049] Reading the device model, version number, and patch number from MIB objects
Before modification: The device does not support reading the device model, version number, and patch number from MIB objects.
After modification: The device supports reading the device model from the hh3cSysProductName MIB object, version number from the hh3cSysProductVersion MIB object, and patch number from the hh3cSysPatchPlatVersion MIB object.
Remarks: In an IRF fabric, only the device model, version number, and patch number of the master device can be retrieved from MIB objects.
- [202301310408] Modified the restrictions on the source port and monitor port in a mirroring group
Before modification: The source port and monitor port in a mirroring group cannot both be aggregate interfaces or aggregation member ports.
After modification: The source port and monitor port in a mirroring group can both be aggregate interfaces or aggregation member ports.
- [202306091868] Assigning an interface to a VLAN on the Web interface
Before modification: When you log in to the device through the Web interface, you cannot assign an interface to a VLAN.
After modification: When you log in to the device through the Web interface, you can assign an interface to a VLAN.

Operation changes in R6710P01

None.

Operation changes in R6710

- [202212070921] Added support for NETCONF/gRPC collection of the system power consumption in real time
Before modification: The chassis and boards data collected through NETCONF/gRPC does not include the total power, residual power, and nominal power data.
After modification: The chassis and boards data collected through NETCONF/gRPC includes the total power, residual power, and nominal power data.
- [202209051705] IPv6 addresses of BGP peers that exceed 15 characters in the display bgp peer command output are displayed in one line
Before modification: IPv6 addresses of BGP peers that exceed 15 characters in the **display bgp peer** command output are displayed in two lines.

After modification: IPv6 addresses of BGP peers that exceed 15 characters in the **display bgp peer** command output are displayed in one line.

- [202208151780] Changed the default state of source MAC address learning for Layer 2 protocol packets

Before modification: By default, the device learns the source MAC addresses of Layer 2 protocol packets (the **mac-address mac-learning pdu** command is used).

After modification: By default, the device is disabled from learning the source MAC addresses of Layer 2 protocol packets (the **undo mac-address mac-learning pdu** command is used).

MAC addresses of the following protocols are involved:

- BPDUs destined for a MAC address in the range of 0x01-80-c2-00-00-00 to 0x01-80-c2-00-00-0f.
- GARP PDUs destined for a MAC address in the range of 0x01-80-c2-00-00-20 to 0x01-80-c2-00-00-2f.
- PVST BPDUs destined for MAC address 0x01-00-0c-cc-cc-cd.

Operation changes in F6705

- [202204201853] Optimized the resource alarm log

Before modification: In the resource alarm log, the Total value equals the sum of the Used value and the Free value.

After modification: In the resource alarm log, the Total value equals the total number of device resources.

- [202112081667] Modified the maximum value to 9008 for the MTU in VSI view and IP MTU in interface view

Before modification: The maximum value is 1500 for the MTU in VSI view and IP MTU in interface view.

After modification: The maximum value is 9008 for the MTU in VSI view and IP MTU in interface view.

- [202204151769] Added support of the CRC error packet alarm function for printing log messages with detailed port numbers

Before modification: For the CRC error packet alarm function, log messages printed do not contain detailed port numbers.

After modification: For the CRC error packet alarm function, log messages printed contain detailed port numbers.

- [202205301866] When the parity-error consistency-check log enable, parity-error consistency-check threshold, parity-error unrecoverable log enable, or parity-error unrecoverable reboot command is executed in system view, the configuration is displayed in the configuration file

Before modification: When the following commands are executed in system view, the configuration is not displayed in the configuration file.

- **parity-error consistency-check log enable**
- **parity-error consistency-check threshold**
- **parity-error unrecoverable log enable**
- **parity-error unrecoverable reboot**

After modification: When the following commands are executed in system view, the configuration is displayed in the configuration file.

- **parity-error consistency-check log enable**
- **parity-error consistency-check threshold**

- **parity-error unrecoverable log enable**
- **parity-error unrecoverable reboot**
- [202112171313] Modified feature: Displaying kernel information upon power cycling of a device

Before modification: After you execute the reboot command to reboot the device, the device can display kernel-related commands. After you power cycle the device, the device cannot display kernel-related commands.

After modification: After you power cycle the device, the device cannot display the following kernel-related commands:

 - **display kernel deadlock**: Displays kernel thread deadlock information.
 - **display kernel exception**: Displays kernel thread exception information.
 - **display kernel reboot**: Displays reboot information for the device.
 - **display kernel starvation**: Displays kernel thread starvation information.
- [202201180838] Modified feature: Optimizing downlink interface state switchover time for Monitor Link

Before modification: In a monitor link group, when the uplink interfaces go down, the state switchover for all downlink interfaces takes a relatively long time.

After modification: In a monitor link group, when the uplink interfaces go down, the state switchover for all downlink interfaces takes less than 5 seconds.

Operation changes in E6702

None.

Operation changes in R2717

None.

Operation changes in R2702

- [201810100437] Added support for using existent and nonexistent policies to filter the BGP routes to be advertised to a peer or peer group.
- [201710190032] Added support for MACsec 256-bit encryption to the LSWM18CQMSEC(JH957A) interface module.

Before modification: The LSWM18CQMSEC(JH957A) interface module only supports the default MACsec 128-bit encryption.

After modification: The **macsec cipher-suite gcm-aes-256** and **macsec cipher-suite gcm-aes-128** commands were added. You can use these commands to enable 256-bit encryption and 128-bit encryption on the LSWM18CQMSEC(JH957A) interface module.
- [201811260646] Removed the **I - Invalid** value from the **Flags** field in the output from the **display evpn route nd** command.
- [201708150654] Added DHCP support for OVSDB.
- [201809260394] Added the DRNI MIB.
- [201711140340] Added the EVPN MIB.
- [201811090674] Added the watchdog timer reset time to the last system reboot information.
- [201901090330] Added CPU model information to the output from the **display version** command.

- [201805220468]Added support for configuring OSPF commands in VSI interface view and disabling a VSI interface from receiving and sending OSPF packets.
- [201904130555]Added SNMP and CLI support for reading the current, voltage, fan direction, and power of a 650 W power supply.
- [201804030211]Added 802.1X MAC address information to the output from the **display mac-address interface** command.
- [201804270716]Added NETCONF support for reading packet statistics about VSI interfaces used for EVPN or VXLAN Layer 3 forwarding.
- [201904090672]Added support for inbound and outbound rate limiting for Layer 3 subinterfaces.

Operation changes in F2701

None.

Operation changes in R2612P02

- Added support for using NETCONF to configure the **aggregate** command to create a summary route in the BGP routing table.
Before modification: NETCONF cannot be used to configure the **aggregate** command to create a summary route in the BGP routing table.
After modification: NETCONF can be used to configure the **aggregate** command to create a summary route in the BGP routing table.
- Added support for using NETCONF to configure the AS path attribute of a routing policy.
Before modification: NETCONF cannot be used to configure the AS path attribute of a routing policy.
After modification: NETCONF can be used to configure the AS path attribute of a routing policy.
- Added support for using NETCONF to apply a routing policy to routes of the specified network segment.
Before modification: NETCONF cannot be used to apply a routing policy to routes of the specified network segment.
After modification: NETCONF can be used to configure the **network *ipv4-address* [*mask-length* | *mask*] [**route-policy** *route-policy-name*]** command to apply a routing policy to routes of the specified network segment.
- Added the **Vendor Part Number** field to the **display transceiver interface** command output.
Before modification: The **display transceiver interface** command output does not support the **Vendor Part Number** field.
After modification: The **display transceiver interface** command output supports the **Vendor Part Number** field.
- Added support for using NETCONF to assign a preferred value to routes received from a peer or peer group.
Before modification: NETCONF cannot be used to assign a preferred value to routes received from a peer or peer group.
After modification: NETCONF can be used to configure the **peer { *group-name* | *ipv4-address* [*mask-length*] } **preferred-value** *value*** command to assign a preferred value to routes received from a peer or peer group.
- Modified the bandwidth value range for traffic of an AC.
Before modification: The bandwidth value range for traffic of an AC is 64 to 4194303 kbps.

- After modification: The bandwidth value range for traffic of an AC is 64 to 167772159 kbps.

 - Modified the broadcast, multicast, or unknown unicast restraint bandwidth for a VSI.

Before modification: The broadcast, multicast, or unknown unicast restraint bandwidth value range for a VSI is 0 and 64 to 4194303 kbps.

After modification: The broadcast, multicast, or unknown unicast restraint bandwidth value range for a VSI is 0 and 64 to 167772159 kbps.
 - Modified the maximum value for the CIR and PIR parameters in the **car** command to 4294967288 kpps.

Before modification: The maximum value for the CIR and PIR parameters in the **car** command in traffic behavior view is 160000000 kbps.

After modification: The maximum value for the CIR and PIR parameters in the **car** command in traffic behavior view is 4294967288 kpps.
 - Added support for the coexistence of the MPLS feature and the VXLAN IP gateway feature.

Before modification: The MPLS feature and the VXLAN IP gateway feature are mutually exclusive.

After modification: The MPLS feature and the VXLAN IP gateway feature can coexist.
 - Modified the maximum ARP packet processing rate to 1000 pps.

Before modification: The maximum ARP packet processing rate is 1800 pps.

After modification: The maximum ARP packet processing rate is 1000 pps.
 - No error information is returned when the **tunnel vpn-instance** command is executed in tunnel view.

Before modification: When the **tunnel vpn-instance vpn-instance-name** command is executed in tunnel view, the system prompts that the command is not supported.

After modification: When the **tunnel vpn-instance vpn-instance-name** command is executed in tunnel view, no error information is returned.
 - Added support for using NETCONF to configure the **bestroute igp-metric-ignore** command

Before modification: NETCONF cannot be used to configure the **bestroute igp-metric-ignore** command.

After modification: NETCONF can be used to configure the **bestroute igp-metric-ignore** command.

Operation changes in R2612P01

- Added support for configuring the user-defined device name as **role name-full IP address** in automated VCF fabric deployment

Before modification: When the device is automatically configured through automated VCF fabric underlay deployment, the device name cannot be configured on Director. The device name is fixed in the **role name-last two fields of IP address** format, for example, **leaf-56.101**.

After modification: When the device is automatically configured through automated VCF fabric underlay deployment, the device name can be configured in the **role name-full IP address** format on Director, for example, **leaf-10.10.56.101**.
- By default, the system does not allocate ACL slice resources to IPv6 and allocates ACL slice resources to IPv6 only after IPv6 is configured, and the system does not allocate ACL slice resources to Layer 3 interfaces

Before modification: By default, the system allocates two ACL slice resources to IPv6 and two ACL slice resources to Layer 3 interfaces.

After modification: By default, the system does not allocate ACL slice resources to IPv6 and allocates ACL slice resources to IPv6 only after IPv6 is configured, and the system does not allocate ACL slice resources to Layer 3 interfaces.

Operation changes in R2612

- Changed the value range for the automatically generated RD of an EVPN instance
Before modification: An automatically generated RD is in the *N: VXLAN ID* format. The initial value of *N* is 0.
After modification: An automatically generated RD is in the *N: VXLAN ID* format. The initial value of *N* is 1.
- Added support for packet statistics of VSI interfaces to the MIB
Before modification: The MIB does not provide packet statistics of VSI interfaces.
After modification: The MIB provides packet statistics of VSI interfaces.
- Modified the sequence number assignment rule for controller connections
Before modification: The switch is assigned the same sequence number each time it sets up a connection to the controller.
After modification: The switch is assigned a different sequence number each time it sets up a connection to the controller.
- Added support for holding the connection to the master controller on the same IRF member switch after an IRF master/subordinate switchover
Before modification: After an IRF master/subordinate switchover, an IRF member switch is randomly selected to establish the connection to the master controller.
After modification: After an IRF master/subordinate switchover, the connection to the master controller is held on the same IRF member switch.
- Removed conflicts between the multicast VXLAN tunnel flood proxy feature and the ARP flood suppression feature
Before modification: When both the multicast VXLAN tunnel flood proxy feature and the ARP flood suppression feature are enabled, the switch cannot operate correctly.
After modification: If the **vxlan tunnel flooding-proxy** command is not executed, ARP flood suppression is enabled on VSIs. If the **vxlan tunnel flooding-proxy** command is executed, ARP flood suppression is disabled on VSIs.
- Increased the maximum number of ACs mapped to a VSI
Before modification: A maximum of 256 ACs can be mapped to a VSI.
After modification: A maximum of 1024 ACs can be mapped to a VSI.
- Added support for OVSDB controller-issued static routes with VSI interfaces as outgoing interfaces.
Before modification: The switch does not accept OVSDB controller-issued static routes with VSI interfaces as outgoing interfaces.
After modification: The switch accepts OVSDB controller-issued static routes with VSI interfaces as outgoing interfaces.
- Only one ACL resource is used by multiple users performing MAC authentication
Before modification: When multiple users perform MAC authentication, multiple ACL resources are used.
After modification: When multiple users perform MAC authentication, only one ACL resource is used.
- CPU cores monitor each other
Before modification: CPU cores cannot monitor each other.

- After modification: CPU cores can monitor each other. If a CPU core fails because of deadlock, the device automatically reboots.
- Added support for MAC address configuration for Layer 3 interfaces and Layer 3 aggregate interfaces when the VXLAN hardware resource mode is border

Before modification: In border mode, you cannot configure MAC addresses for Layer 3 interfaces and Layer 3 aggregate interfaces.

After modification: In border mode, you can configure MAC addresses for Layer 3 interfaces and Layer 3 aggregate interfaces.
 - Added support for packet statistics on Ethernet subinterfaces and Layer 3 aggregate subinterfaces

Before modification: Ethernet subinterfaces and Layer 3 aggregate subinterfaces do not support the **traffic-statistic enable** command.

After modification: Ethernet subinterfaces and Layer 3 aggregate subinterfaces support the **traffic-statistic enable** command.
 - Changed the way of processing IP packets destined for class E addresses.

Before modification: The switching chips forward IP packets destined for class E addresses.

After modification: The switching chips drop IP packets destined for class E addresses.
 - Added 802.1X MAC address information to the output from the **display mac-address interface** command.

Before modification: The **display mac-address interface** command does not display 802.1X MAC address information.

After modification: The **display mac-address interface** command displays 802.1X MAC address information.
 - Changed the log message printed when DRNI configuration consistency check is finished after the IPL comes up.

Before modification: The log message is DRNI/6/DRNI_GLBCONSISTENCYCHECK_FAILURE: Global type XXX configuration consistency check failed.

After modification: The log message is DRNI/6/DRNI_IFCHECK_CONSISTENCY: Finished DR interface XXX type XXX configuration consistency check. No inconsistency exists.
 - Support for generating a log when a parity error occurs on an MMU hardware entry was added.

Before modification: No log is generated when a parity error occurs on an MMU hardware entry.

After modification: Log message "MMU checking error is detected, rebooting the slot is recommended" is generated when a parity error occurs on an MMU hardware entry.
 - Sending packets without VLAN tags if the VLAN corresponding to an AC is the PVID of the interface when VLAN-based VXLAN assignment is configured

Before modification: When the VLAN corresponding to an AC is the PVID of the interface, packets are sent with VLAN tags.

After modification: When the VLAN corresponding to an AC is the PVID of the interface, packets are sent without VLAN tags.

Operation changes in R2610

- Added VCF controller support for issuing DRNI configuration

Before modification: VCF controllers do not support issuing DRNI configuration.

After modification: VCF controllers can issue DRNI configuration.

Operation changes in R2609

- Support for generating a log when a parity error occurs on an MMU hardware entry was added.
Before modification: No log is generated when a parity error occurs on an MMU hardware entry.
After modification: Log message "MMU checking error is detected, rebooting the slot is recommended" is generated when a parity error occurs on an MMU hardware entry.
- Error code "resource-not-ready" was added to NETCONF.
Before modification: Error message "The VRF instance is being deleted." is prompted if NETCONF attempts to issue the operation of adding a VRF to the switch immediately after NETCONF successfully issues the operation of deleting the VRF, because the switch is still deleting the VRF.
After modification: Error code "resource-not-ready" is returned if NETCONF attempts to issue the operation of adding a VRF to the switch immediately after NETCONF successfully issues the operation of deleting the VRF, because the switch is still deleting the VRF.
- Layer 3 Ethernet interface used as a public network interface supports tagging outgoing packets with VLAN 0.
Before modification: On a VPLS network, when a Layer 3 Ethernet interface is used as a public network interface, it tags outgoing packets with VLAN 4095.
After modification: On a VPLS network, when a Layer 3 Ethernet interface is used as a public network interface, it tags outgoing packets with VLAN 0.
- Support for counting traffic both in packets and in bytes after a CAR action is configured in a traffic behavior was added.
Before modification: After a CAR action is configured in a traffic behavior of a QoS policy, only the number of packets or the number of bytes is displayed.
After modification: After a CAR action is configured in a traffic behavior of a QoS policy, both the number of packets and the number of bytes are displayed.
- Changed the prompt for MAC authentication failure in open authentication mode.
Before modification: In open authentication mode, if a user fails MAC authentication, the system prompts for MAC authentication success and allows the user to come online.
After modification: In open authentication mode, if a user fails MAC authentication, the system prompts for MAC authentication failure and allows the user to come online.

Operation changes in R2608P02

- Added support for the rewrite inbound tag nest command on Ethernet service instances of an interface configured with QinQ
Before modification: When an interface is configured with QinQ, Ethernet service instances on the interface cannot be configured with the **rewrite inbound tag nest** command.
After modification: When an interface is configured with QinQ, Ethernet service instances on the interface can be configured with the **rewrite inbound tag nest** command.
- Modified the FreeRatio calculation method in the display memory summary command output
Before modification: In the **display memory summary** command output, the cache is excluded from the FreeRatio calculation.
After modification: In the **display memory summary** command output, the cache is included in the FreeRatio calculation.
- Removed the buildrun for the VSI interfaces deployed by the controller
Before modification: Buildrun is performed for the VSI interfaces deployed by the controller.
After modification: Buildrun is not performed for the VSI interfaces deployed by the controller.

- Increased the maximum WFQ weight to 127
Before modification: The maximum WFQ weight is 15.
After modification: The maximum WFQ weight is 127.
- Increased the number of loopback interfaces to 1K
Before modification: The number of loopback interfaces is 128.
After modification: The number of loopback interfaces is 1K.
- Modified the GRE encapsulation port number to 0x88be for Layer 3 remote port mirroring
Before modification: The GRE encapsulation port number for Layer 3 remote port mirroring is 0x6558.
After modification: The GRE encapsulation port number for Layer 3 remote port mirroring is 0x88be.
- Modified the ND entry specifications in hardware resource modes 1, 2, and 3
Before modification:
 - In hardware resource mode 1, the ND entry specification is 28K.
 - In hardware resource mode 2, the ND entry specification is 44K.
 - In hardware resource mode 3, the ND entry specification is 60K.
 After modification:
 - In hardware resource mode 1, the ND entry specification is 40K.
 - In hardware resource mode 2, the ND entry specification is 72K.
 - In hardware resource mode 3, the ND entry specification is 104K.
- Modified the minimum Ethernet OAM handshake packet transmission interval from 500 ms to 100 ms and the minimum Ethernet OAM connection timeout timer from 1000 ms to 200 ms
Before modification: The minimum Ethernet OAM handshake packet transmission interval is 500 ms and the minimum Ethernet OAM connection timeout timer is 1000 ms.
After modification: The minimum Ethernet OAM handshake packet transmission interval is 100 ms and the minimum Ethernet OAM connection timeout timer is 200 ms.
- Increased the maximum number of BFD sessions from 32 to 64
Before modification: The maximum number of BFD sessions is 32.
After modification: The maximum number of BFD sessions is 64.
- Removed support of VXLAN tunnel interfaces for TTL, DF bit, and ToS configurations
Before modification: The **tunnel ttl**, **tunnel dfbit enable**, and **tunnel tos** commands are supported in VXLAN tunnel interface view.
After modification: The **tunnel ttl**, **tunnel dfbit enable**, and **tunnel tos** commands are supported in VXLAN tunnel interface view.
- Removed support for issuing the ACL flow entry set_filed ip_dst
Before modification: The device does not support the ACL flow entry set_filed ip_dst. However, the rule can be issued, but it does not take effect.
After modification: The device does not support the ACL flow entry set_filed ip_dst. The rule fails to be issued.
- Added support for the mac-address mac-learning priority high and mac-address max-mac-count commands on interfaces on PEXs in an IRF3.1 system
Before modification: Interfaces on PEXs in an IRF3.1 system do not support the **mac-address mac-learning priority high** or **mac-address max-mac-count** command.
After modification: Interfaces on PEXs in an IRF3.1 system support the **mac-address mac-learning priority high** and **mac-address max-mac-count** commands.
- Added the check for switching chip DMA and switching logic components

This software version added the check for switching chip DMA and switching logic components to determine whether they are running correctly.

- Add the support for configuring static ARP entries through NETCONF.
Before modification: Configuring static ARP entries through NETCONF is not supported.
After modification: Configuring static ARP entries through NETCONF is supported.
- Added conflict detection when an AC receives an ARP packet sourced from a gateway address
Before modification: When distributed VXLAN IP gateways exist, no conflict prompt is displayed if an AC receives an ARP packet sourced from a gateway address.
After modification: When distributed VXLAN IP gateways exist, if an AC receives an ARP packet sourced from a gateway address, the switch performs conflict detection and sends a gratuitous ARP response.
- Changed the action to take on an ARP packet received on an AC if the packet is sourced from a gateway address
Before modification: If proxy ARP is enabled, and the switch does not have ARP entries for distributed VXLAN IP gateways, ARP packets sourced from a gateway address are broadcast.
After modification: ARP packets sourced from a gateway address are terminated.
- Removed STP state limitation on traffic forwarding of a link aggregation management port
Before modification: The management port of a link aggregation group does not forward traffic of management VLANs if its STP state is not forwarding.
After modification: The STP state does not affect traffic forwarding of the management port of a link aggregation group.

Operation changes in F2608

- Added VXLAN tunnels' support for sharing a destination IP address.
Before modification: On a VTEP, VXLAN tunnels must use different destination IP addresses.
After modification: On a VTEP, multiple VXLAN tunnels can share a destination IP address.
- Distinguished user-configured MAC addresses of Layer 3 interfaces and subinterfaces by VLAN.
Before modification: MAC addresses manually assigned to Layer 3 interfaces and subinterfaces are not distinguished by VLAN.
After modification: MAC addresses manually assigned to Layer 3 interfaces and subinterfaces are distinguished by VLAN.
- Removed the VXLAN hardware resource mode and the **switch-mode** setting from automatically issued configurations for leaf nodes.
Before modification: The VXLAN hardware resource mode and the **switch-mode** setting are directly issued to leaf nodes.
After modification: The VXLAN hardware resource mode and the **switch-mode** setting are automatically set based on configurations and network environment.
- Modified the maximum number of VSANs supported by the switch from 16 to 32.
- Added support for equal cost LSPs when the switch acts as a transit node in an MPLS network.
Before modification: The switch does not support equal cost LSPs when acting as a transit node.
After modification: The switch supports equal cost LSPs when acting as a transit node.
- Modified the prompt displayed when an interface of a PEX is configured as a mirroring source.
Before modification: When the **mirroring-group group-id mirroring-port inbound** command is executed on an interface of a PEX, **Mirroring setting failed.** is displayed.

After modification: When the **mirroring-group group-id mirroring-port inbound** command is executed on an interface of a PEX, **The operation is not supported.** is displayed.

- Added support for the **speed** command on breakout interfaces.
Before modification: The **speed** command is not supported in breakout interface view.
After modification: The **speed** command is supported in breakout interface view.
- Added support for ECMP to VXLAN multicast forwarding.
Before modification: VXLAN multicast packets cannot be forwarded through ECMP routes.
After modification: VXLAN multicast packets can be forwarded through ECMP routes.
- Added NETCONF support for configuring MPLS.
Before modification: NETCONF does not support configuring MPLS.
After modification: NETCONF supports configuring MPLS.
- Added CLI and NETCONF support for configuring descriptions for traffic classes in QoS policies.
Before modification: Descriptions cannot be configured for traffic classes in QoS policies through the CLI or NETCONF.
After modification: Descriptions can be configured for traffic classes in QoS policies through the CLI and NETCONF.
- Added BFD support for VXLAN flood proxy tunnels.
Before modification: BFD does not support detecting connectivity of VXLAN flood proxy tunnels.
After modification: BFD supports detecting connectivity of VXLAN flood proxy tunnels.
- Modified candidate upstream member interface requirements for IRF 3.1 when ports of extension interface cards are used to connect PEXs to the parent fabric.
Before modification: For 5940 2-slot Switch JH397A and 5940 4-slot Switch JH398A switches, you can use only the lowest numbered two QSFP+ ports on the 5930 8-port QSFP+ Module JH183A card installed in slot 1. For the 5940 32QSFP+ JH396A switch, you can use the highest numbered four QSFP+ ports on the switch.
After modification: You can use ports of an extension interface card installed in any slot.
 - For 5940 2-slot Switch JH397A, 5940 4-slot Switch JH398A switches, you can use the following ports:
 - Highest numbered two QSFP+ ports on an 5930 8-port QSFP+ Module JH183A card.
 - Highest numbered six SFP+ ports and highest numbered two QSFP+ ports on an 5930 24p 10GBase-T and 2p QSFP+ Mod JH182A, 5930 24p SFP+ and 2p QSFP+ Mod w Msec JH181A, 5930 24p SFP+ and 2p QSFP+ Mod JH180A or 5930 24p CP and 2p QSFP+ Mod JH184A card.
 - For the 5940 32QSFP+ JH396A switch, you can use the highest number eight QSFP+ ports on the switch.
- Modified the maximum number of IGMP snooping entries from 8 K to 16 K when the **switch-mode** setting is set to 1, 2, or 3.
- Modified the maximum number of isolation groups from 8 to 128.
- Added support for management Ethernet interface backup for IRF.
Before modification: IRF member devices do not support management Ethernet interface backup.
After modification: When the management Ethernet interface of the master goes down, the management Ethernet interface on a subordinate is activated to take over.
- Optimized OSPF neighboring relationship establishment (applicable only to switches on Tencent networks)

Before modification: If a receiving LSA is older than or the same as the local LSA and is being requested, the neighbor state is changed to **Exstart**. The switch initiates a master negotiation and does not establish neighboring relationship with the neighbor.

After modification: If a receiving LSA is older than or the same as the local LSA and is being requested and to be flooded, the neighbor state is not changed to **Exstart**. The switch establishes neighboring relationship with the neighbor.

Operation changes in F2606

- OpenFlow connections will not be disconnected because of short link flapping.
Before modification: OpenFlow connections will be disconnected because of short link flapping.
After modification: OpenFlow connections will not be disconnected because of short link flapping.

Operation changes in F2605

- Priority of a QoS policy applied to an Ethernet service instance became higher than that applied to an interface
Before modification: When both an interface and an Ethernet service instance on the interface have QoS policies applied, if a packet matches both QoS policies, the QoS policy applied to the interface takes effect.
After modification: When both an interface and an Ethernet service instance on the interface have QoS policies applied, if a packet matches both QoS policies, the QoS policy applied to the Ethernet service instance takes effect.
- Response to the Ctrl+C interruption operation during the Python script execution process
Before modification: After the device is started, the device does not respond to the Ctrl+C interruption operation during the Python script execution process.
After modification: After the device is started, the device responds to the Ctrl+C interruption operation during the Python script execution process.
- VLAN interfaces do not forward packets received from CPU interfaces to other Layer 3 Ethernet subinterfaces
Before modification: ARP and OSPF packets sent out of a VLAN interface are copied and sent out of Layer 3 Ethernet subinterfaces.
After modification: ARP and OSPF packets sent out of a VLAN interface are not copied and sent out of Layer 3 Ethernet subinterfaces.
- Added process monitoring and maintenance commands to the probe view
Before modification: No process monitoring and maintenance commands are available in probe view.
After modification: The **process restart**, **process start**, and **process shutdown** commands for process monitoring and maintenance are added to the probe view.
- Buildrun when the **hardware-resource** command uses the default setting
Before modification: When the **hardware-resource** command uses the default setting, buildrun is not performed, and pre-reading and issuing cannot be configured.
After modification: When the **hardware-resource** command uses the default setting, buildrun can be performed, and pre-reading and issuing can be configured.
- Added three frame match criteria to Ethernet service instances on PEXs when VXLAN is configured in an IRF3.1 system
When VXLAN is configured in an IRF3.1 system and the AC is an Ethernet service instance on a PEX, the following match criterion configuration commands are added to the Ethernet service

instance: **encapsulation s-vid** *vlan-id-list*, **encapsulation default**, and **encapsulation tagged**.

- Modified the maximum number of next hops for tunnels that share a physical outgoing interface.
Before modification: Tunnels that share a physical outgoing interface can have only one next hop.
After modification: Tunnels that share a physical outgoing interface can have multiple next hops.

Operation changes in F2604

- Added support for sharing of MAC addresses among Layer 3 interfaces
Before modification: Two Layer 3 interfaces cannot use the same MAC address.
After modification: Two Layer 3 interfaces can share a MAC address.
- Added the log message for applying a traffic redirecting QoS policy to a Layer 3 aggregate interface
Before modification: If a QoS policy configured with traffic redirecting is applied to a Layer 3 aggregate interface, the configuration fails and no prompt is displayed.
After modification: If a QoS policy configured with traffic redirecting is applied to a Layer 3 aggregate interface, the message "The operation is not supported" is displayed.
- Changed the maximum number of ARP entries
Before modification: In the output from the **display arp entry-limit** command, the maximum number of ARP entries is 8192.
After modification: In the output from the **display arp entry-limit** command, the maximum number of ARP entries is 32768.
- Changed the resolution method for short static ARP entries configured for Layer 3 Ethernet subinterfaces
Before modification: If a short static ARP entry is configured for a Layer 3 Ethernet subinterface, the entry is resolved after the switch sends an ARP request whose target IP address is the IP address in the short entry and receives a reply.
After modification: If a short static ARP entry is configured for a Layer 3 Ethernet subinterface enabled with unambiguous VLAN termination, the entry is resolved after it is configured. If ambiguous VLAN termination is configured for the Layer 3 Ethernet subinterface, the entry is resolved after the switch sends an ARP request whose target IP address is the IP address in the short entry and receives a reply.
- Removed MACsec support for Layer 3 Ethernet interfaces, member ports of service loopback groups, and IRF physical interfaces
Before modification: MACsec can be enabled on Layer 3 Ethernet interfaces, member ports of service loopback groups, and IRF physical interfaces.
After modification: MACsec cannot be enabled on Layer 3 Ethernet interfaces, member ports of service loopback groups, and IRF physical interfaces. If you enable MACsec on these interfaces, the system displays a prompt.
- Added the function of synchronizing OVSDB QoS configuration to the **vtep-ctl** command in probe view
Before modification: The **vtep-ctl** command does not support synchronizing OVSDB QoS configuration.
After modification: The **vtep-ctl** command supports synchronizing OVSDB QoS configuration.
- Added the support for multiuser login of VFC interfaces
Before modification: A VFC interface does not allow multiple users to log in to the device concurrently.

- After modification: A VFC interface allows multiple users to log in to the device concurrently.
- Added the radar detection and link emulation features

Before modification: The device does not provide a NETCONF interface for the Director to send detection packets.

After modification: The device provides a NETCONF interface for the Director to send detection packets and the device sends log messages to the controller when receiving detection packets.
- Added dynamic ACL assignment to MAC authentication

Before modification: MAC authentication does not support dynamic ACL assignment.

After modification: MAC authentication supports dynamic ACL assignment.
- Added the feature of collecting packets statistics for aggregate subinterfaces and Layer 3 Ethernet subinterfaces through MIB

Before modification: This feature is not supported.

After modification: This feature is supported.
- Modified the Ethernet frame match criteria when the controller deploys the VLAN=0 match criterion.

Before modification: When the controller deploys the VLAN=0 match criterion, the Ethernet frame match criterion is **encapsulation default**.

After modification: When the controller deploys the VLAN=0 match criterion, the Ethernet frame match criterion is **encapsulation untagged**.

Operation changes in F2603

- For the 5940 48XGT 6QSFP28 JH391A/5940 48SFP+ 6QSFP28 JH390A device, added support for using 100-GE ports that operate at 40 Gbps as IRF physical interfaces.

Before modification: 100-GE ports operating at 40 Gbps cannot be used as IRF physical interfaces.

After modification: 100-GE ports operating at 40 Gbps can be used as IRF physical interfaces.
- Modified the mode for forwarding an ARP broadcast packet received on an AC or tunnel interface.

Before modification: When ARP flood suppression is enabled and an ARP broadcast packet arrives at an AC or tunnel interface, the software replicates a copy for each AC and tunnel interface. Then, the software forwards the copies to the ACs and tunnel interfaces.

After modification: When ARP flood suppression is enabled and an ARP broadcast packet arrives at an AC or tunnel interface, the software forwards only one copy. Then, the hardware replicates a copy for each AC and tunnel interface and forwards the copies to the ACs and tunnel interfaces.
- Changed the default MTU of VSI interfaces from 1500 bytes to 1444 bytes.
- Added support for connecting member ports of two local Layer 3 dynamic aggregate interfaces.
 - Before modification: If two Layer 3 Ethernet interfaces on the device are assigned to different dynamic aggregation groups, the interfaces cannot be Selected when they are connected.
 - After modification: If two Layer 3 Ethernet interfaces on the device are assigned to different dynamic aggregation groups, the interfaces can be Selected when they are connected.
- Optimized ACL resource allocation for service chain flow entries.
 - Before modification: Each flow entry is assigned one ACL resource.
 - After modification: Flow entries with the same service ID and tunnel interface index are assigned on ACL resource.
- Added support for redistributing IGP routes to the BGP routing table.

- Before modification: IGP route redistribution is not supported for the BGP routing table.
- After modification: The **import-route protocol** [{ *process-id* | **all-processes** } [**allow-direct** | **med med-value** | **route-policy route-policy-name**] *] command was added to redistribute IGP routes to the BGP routing table.
- Added Neutron support for OpenStack Mitaka.
- Added traffic statistic collection for VSI interfaces, VLAN interfaces, and Layer 3 aggregate subinterfaces.
- Added NETCONF support for import and export routing policy configuration of BGP EVPN peers.
- Added NETCONF support for the **igmp-snooping enable** command in VSI view.
- Changed the default value of the IRF bridge MAC persistence timer from 6 minutes to 12 minutes.
- Changed the maximum number of routes a VPN instance can have in hardware resource mode 4.
 - Before modification: A VPN instance can have a maximum of 128 K routes.
 - After modification: The maximum number of routes that a VPN instance supports equals the maximum number of routes that the switch supports.

Operation changes in R2509P02

- Modified the default broadcast restraint bandwidth, multicast restraint bandwidth, and unknown unicast restraint bandwidth.
 - Before modification: The default broadcast restraint bandwidth, multicast restraint bandwidth, and unknown unicast restraint bandwidth are 5120 kbps.
 - After modification: The default broadcast restraint bandwidth, multicast restraint bandwidth, and unknown unicast restraint bandwidth are not limited.
- Modified the default broadcast, multicast, and unknown unicast suppression bandwidth for VSIs

Before modification: The default suppression bandwidth is 5120 kbps.

After modification: The default suppression bandwidth is 4096 kbps.
- Modified the VLAN check function when a MAC address is set for a Layer 3 Ethernet interface or VLAN interface.

Before modification: The system checks whether the newly set MAC address belongs to the same VLAN as an existing entry and issues different entries for the same MAC address and different VLANs.

After modification: The system does not check whether the newly set MAC address belongs to the same VLAN as an existing entry and issues only one entry for one MAC address.
- Modified the **Input Interface** field for displaying outgoing unicast packets sampled by using sFlow

Before modification: This field displays **N/A**.

After modification: This field displays the input interface name.

Operation changes in R2509P01

None.

Operation changes in R2509

- Added log messages for parity error recovery failures.
Before modification: The switch does not generate log messages for parity error recovery failures.
After modification: The switch generates a log message when a parity error recovery failure occurs.
- Increased the maximum number of VSI interfaces from 99999 to 16777215.
- Added support for the **hardware-resource routing-mode ipv6-128** command when the **switch-mode 4** command is executed.
- Increased the IRF master election time from 20 seconds to 40 seconds.
- Added support for traffic statistics of VSI interfaces, VLAN interfaces, and Layer 3 aggregate subinterfaces.

Operation changes in R2508

- Changing the maximum rate allowed by software and hardware rate limiting in PVST
Before modification: The maximum rate allowed is 200 pps for software rate limiting and 768 pps for hardware rate limiting.
After modification: The maximum rate allowed is 2560 pps for software rate limiting and 1536 pps for hardware rate limiting.

Operation changes in R2507

None.

Operation changes in R2506

First release.

Restrictions and cautions

Restrictions

- If you execute the flow-control hardware command in console line view or console line class view, the console port of the device will not respond.

Open problems and workarounds

201612070474

- Symptom: MACsec is enabled on interfaces of an interface card. When the interface card is removed and re-installed, flapping occurs on other interfaces of the switch that have established MACsec connections.
- Condition: This symptom might occur if the following conditions exist:
 - a. MACsec is enabled on interfaces of an interface card and other interfaces on the switch.

- b. The interface card is removed and re-installed.
- Workaround: Use the following procedure to re-install the interface card:
 - c. Delete all MACsec settings on the interface card.
 - d. Remove the interface card and re-install it.
 - e. Reconfigure MACsec on the interface card.

202012310742

- Symptom: An IRF fabric fails to forward traffic for an AC that matches untagged traffic.
- Condition: This symptom might occur if the following operations are performed:
 - a. Create an aggregate interface with member ports on multiple IRF member devices.
 - b. Configure an AC to match untagged traffic on the aggregate interface.
 - c. Reboot an IRF member device.
 - d. Restart the aggregate interface, or delete another AC that matches untagged traffic and then recreate it.
 - e. Execute the **flooding disable** command on the VSI mapped to the AC created in step b.
- Workaround: Do not execute the **flooding disable** command.

202302200107

- Symptom: Traffic forwarding is interrupted after a traffic outgoing interface switchover occurs on an DRNI system.
- Condition: This symptom might occur if an DRNI member device receives ND entries from the DRNI peer after a device restart and a traffic outgoing interface switchover occurs on the device.
- Workaround: If this problem has occurred, delete the ND entries for the DRNI member devices to learn them again.

202307130980

- Symptom: ARP and ND entries of DRNI extra VLANs cannot be synchronized over the peer link.
- Condition: This symptom might occur if an DRNI member device reboots or its peer-link interface flaps.
- Workaround: None.

List of resolved problems

Resolved problems in R6710P03

202305040706

- Symptom: An interface on an 5940 switch comes up and goes down repeatedly after a 10G transceiver module is inserted into the interface with a QSFP+ to SFP+ adapter and no optical fiber is connected to the transceiver module
- Condition: This symptom occurs if a 10G transceiver module is inserted into an interface with a QSFP+ to SFP+ adapter on an 5940 switch and no optical fiber is connected to the transceiver module.
- Remarks: None.

202302141672

- Symptom: gRPC-sampled data reported is abnormal irregularly.

- Condition: This symptom occurs if you configure gRPC sensor path **ifmgr/trafficstatistics/interfaces**.
- Remarks: None.

202302150003

- Symptom: The log file **fabric.log** generated by VCF fabric exhausts the memory.
- Condition: This symptom occurs if the automated deployment scenario of VCF fabric runs for a long period of time or interfaces flap.
- Remarks: None.

202302151795

- Symptom: When member ports of an aggregate interface come up, the member ports will flap once and then restore to up.
- Condition: This symptom occurs if the **lACP period short** and **link-delay down 5** commands are executed on the aggregation member ports and then wavelength division switchover is performed.
- Remarks: None.

202302141663

- Symptom: Dynamic ND entries do not age out on an EVPN DRNI member device.
- Condition: This symptom occurs if an EVPN DRNI member device reboots.
- Remarks: None.

202206170581

- Symptom: The **mka enable** command is executed to enable MACsec on an LSWM18CQMSEC module, and then the module is rebooted by using the **reboot** command or is reinstalled. After the module reboots, the **mka enable** setting configured in interface view is lost.
- Condition: This symptom might occur if the following operations are performed:
 - a. Execute the **mka enable** command to enable MACsec on an LSWM18CQMSEC module.
 - b. Reboot the module with the **reboot** command or reinstall the module.
- Remarks: None.

202306080897

- Symptom: The device generates message **Failed to save license data to the primary license storage area** at intervals of 24 hours.
- Condition: This symptom occurs when the system fails to read and write the license storage area because of flash memory failure.
- Remarks: None.

202305291927

- Symptom: API Device/Base cannot be read on Postman.
- Condition: This symptom might occur when you use Postman to retrieve the Device/Base node.
- Remarks: None.

202305100224

- Symptom: Protocol packets are dropped in an EVPN VXLAN-DCI network.
- Condition: This symptom occurs if the TTL of the protocol packets is 1.
- Remarks: None.

202211031872

- Symptom: During the ISSU loading process, one IRF member device experiences packet loss for approximately 18 seconds.
- Condition: This symptom occurs if EVPN VXLAN is configured on IRF member devices, a subordinate member device is restarted, and Layer 3 VXLAN traffic by default matches a blackhole route.
- Remarks: None.

202303131038

- Symptom: In the output from the display ipv6 interface command, the IPv6 address, interface name, and VPN fields are displayed on different lines, which should be displayed on the same line.
- Condition: This symptom occurs if you execute the display ipv6 interface command.
- Remarks: None.

202305221758

- Symptom: When the outgoing interface of a VXLAN tunnel is a Layer 3 aggregate interface, the outgoing VXLAN packets carry VLAN tag 4095 unexpectedly. As a result, the peer cannot learn ARP entries.
- Condition: This symptom occurs if the outgoing interface of a VXLAN tunnel is a Layer 3 aggregate interface.
- Remarks: None.

202305300007

- Symptom: Creation of a VSI interface, Layer 3 subinterface, or Layer 3 aggregate subinterface might fail.
- Condition: This symptom occurs if a VSI interface, Layer 3 subinterface, or Layer 3 aggregate subinterface is created.
- Remarks: None.

202304240579

- Symptom: Isolation of aggregation member ports no longer takes effect on a DR interface, and the traffic is forwarded between the aggregation member ports.
- Condition: This symptom occurs if the following operations are performed:
 - a. Shut down all aggregation member ports of the IPP and DR interfaces, save the configuration, and reboot the device.
 - b. Bring up the aggregation member ports of the IPP.
 - c. After half of the DRNI restoration delay elapses, bring up the aggregation member ports of the DR interfaces.
- Remarks: None.

202209230460

- Symptom: In gRPC dial-in mode, some sampling paths cannot collect data and the data is collected by other sampling paths.
- Condition: This symptom might occur if you configure multiple sampling paths in gRPC dial-in mode.
- Remarks: None.

202211140499

- Symptom: OSPF BFD flaps repeatedly.
- Condition: This symptom occurs if you use borrowed loopback interface addresses to establish OSPF neighbor relationship, configure BFD for OSPF, and then reboot the device.
- Remarks: None.

202302150003

- Symptom: The log file **fabric.log** generated by VCF fabric exhausts the memory.
- Condition: This symptom occurs if the automated deployment scenario of VCF fabric runs for a long period of time or interfaces flap.
- Remarks: None.

202305110216

- Symptom: On a multicast VXLAN network, multicast traffic cannot be forwarded.
- Condition: This symptom occurs if the device starts with the factory defaults and then you configure multicast VXLAN in the following order: first configure tunnels and VSIs, and then configure multicast.
- Remarks: None.

202304171574

- Symptom: The switch cannot obtain an IPv6 address after it is rebooted, and IPv6 automatic deployment fails.
- Condition: This symptom occurs if the controller deploys the configuration to change the hardware resource mode during automatic deployment and the controller does not assign a fixed IPv6 address.
- Remarks: None.

202305081426

- Symptom: In an EVPN or VXLAN distributed gateway network, when the device receives a tunneled packet with a source IP address the same as a VSI interface address, the device will reply with a gratuitous ARP response, which can lead to high CPU usage.
- Condition: This symptom might occur if the distributed gateways perform ARP probing in response to traffic.
- Remarks: None.

202306100168

- Symptom: A device attached to an DRNI system with dual-active VLAN gateways configured cannot learn ARP information about a peer.
- Condition: This symptom occurs if a device attached to an DRNI system with dual-active VLAN gateways sends an ARP request to obtain the ARP information about a peer.
- Remarks: None.

202305120926

- Symptom: The device gets stuck after a controller deploys the default action to interfaces on the device.
- Condition: This symptom occurs if the device has port security settings and the controller uses multiple sessions to deploy the default action.
- Remarks: None.

202304250098

- Symptom: After the **peer advertise vpn-reoriginate ibgp** command is executed, the local device removes private AS numbers (in the range of 65512 to 65534) from routes before advertising those routes to the specified peers. This operation affects the results of optimal route selection on the peers. When you execute the **display bgp update-group l2vpn evpn** command to view the update group information for the specified peers, the command output displays **Public-AS-Only: Yes**.
- Condition: This symptom occurs if you execute the **peer advertise vpn-reoriginate ibgp** command. This command enables the device to remove private AS numbers (in the range of 65512 to 65534) from routes before the device advertises those routes to the specified peers.
- Remarks: None.

202305100217

- Symptom: When an endpoint sends an RARP message, the route used for forwarding traffic to the endpoint flaps, and traffic loss occurs.
- Condition: This symptom occurs if an endpoint dualhomed or singlhomed to an EVPN DRNI system sends an RARP packet.
- Remarks: None.

202303160020

- Symptom: When a DHCP user comes online, the DHCP process is closed abnormally.
- Condition: This symptom might occur if the following conditions exist:
 - a. The DHCP user comes online through interface 1 and two IP addresses (for example, IP address A and IP address B) are obtained.
 - b. The DHCP user later comes online through interface 2 and IP address A is obtained.
 - c. The clientinfo entries on the DHCP relay device are reset.

202306060566

- Symptom: After OSPF establishes a neighbor relationship with a neighboring device, the neighbor cannot learn the default route advertised by the local device.
- Condition: This symptom might occur if you create OSPF view without associating any interfaces and then execute the **nssa default-route-advertise** command.
- Remarks: None.

202305200093

- Symptom: The device is disconnected from the controller when a patch is installed from the controller.
- Condition: This symptom occurs if you install a patch from the controller and restart the xmlcfd process when the patch is installed.
- Remarks: None.

202206071105

- Symptom: When you configure an **s-vid** (outer VLAN IDs) match criterion for a VPLS Ethernet service instance, you can only specify a single VLAN ID and cannot specify a VLAN ID range.
- condition: This symptom occurs when you configure a packet match criterion for an Ethernet service instance of a VPLS network.

202305220011

- Symptom: IP address conflicts occur between four leaf devices because of inconsistent ARP and MAC information, and the CPU usage of the leaf devices reaches 70%.
- Condition: This symptom occurs if the following conditions exist:

- With ARP proxy enabled, a probe packet is sent when a remote ARP rule for EVPN is withdrawn.
- A probe packet is sent if a remote ARP rule overwrites a local ARP entry.
- Remarks: None.

Resolved problems in R6710P01

202301120578

- Symptom: After an incremental patch is uninstalled, the **display boot-loader** command does not display information about a non-incremental patch.
- Condition: This symptom occurs if both an incremental patch and a non-incremental patch are installed.
- Remarks: None.

Resolved problems in R6710

202208040005

- Symptom: In an M-LAG system that use VLAN interfaces to act as dual-active gateways for the same VLAN, an M-LAG member device cannot successfully ping an endpoint.
- Condition: This symptom occurs if the other M-LAG member device receives ICMP replay packets.
- Remarks: None.

202208241285

- Symptom: A QoS policy applied to a control plane cannot filter the protocol packets to the control plane
- Condition: This symptom occurs when you apply a QoS policy to a control plane to filter protocol packets.
- Remarks: None.

202211010383

- Symptom: When a client-oriented MACsec connection is established between an Aruba device and HPE switch, the MACsec protocol cannot come up, and the connection cannot be established correctly.
- Condition: This symptom occurs if a client-oriented MACsec connection is established between an Aruba device and HPE switch.
- Remarks: None.

202204071026

- Symptom: A QoS policy applied to a VSI takes effect only on traffic forwarded at Layer 2 and does not take effect on traffic forwarded at Layer 3.
- Condition: This symptom occurs if a QoS policy is applied to a VSI.
- Remarks: None.

202211031872

- Symptom: On an EVPN DCI network, the device forwards multiple copies of a packet.
- Condition: This symptom occurs if the automatic MAC address learning feature is not enabled on the tunnel side, and the Layer 3 service traffic received on the tunnel side does not match any MAC address entry.

- Remarks: None.

202211050218

- Symptom: After the BFD MAD configuration is deleted from a VLAN interface, the configuration remains.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure BFD MAD on the VLAN interface, and bind the VLAN interface to a VPN instance.
 - b. Configure BFD MAD on an aggregate interface. Bind the aggregate interface to the same VPN instance as the VLAN interface.
 - c. Delete the BFD MAD configuration from the VLAN interface.
 - d. Delete the VLAN interface configured with BFD MAD.
- Remarks: None.

202211050189

- Symptom: After an IRF member device is rebooted, the **display bfd session** command output displays two BFD MAD sessions.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure BFD MAD on an aggregate interface, and bind the aggregate interface to a VPN instance.
 - b. Delete the BFD MAD configuration from the aggregate interface.
 - c. Configure BFD MAD on a VLAN interface. Bind the VLAN interface to the same VPN instance as the aggregate interface.
 - d. Configure BFD MAD on the aggregate interface again.
 - e. Reboot an IRF member device. After the device is rebooted and the IRF fabric is formed again, execute the **display bfd session** command to display the BFD MAD sessions.
- Remarks: None.

202204010582

- Symptom: On an EVPN network, some aggregation member ports do not forward traffic.
- Condition: This symptom occurs if the following conditions exist:
 - An IRF fabric acts as a VTEP.
 - A VXLAN tunnel has multiple equal-cost next hops, and the next hops correspond to multi-card aggregate interfaces.
 - The local aggregation member ports are shut down and then brought up.
- Remarks: None.

202204110427

- Symptom: On a VPLS network with primary/backup PWs, traffic might fail to be forwarded after primary/backup PW switchover.
- Condition: This symptom occurs if the public network interface of the primary PW flaps to switch traffic between the primary and backup PWs.
- Remarks: None.

202204090439

- Symptom: The console gets stuck after repeated execution of the **port-security enable** or **port-security port-mode** command.
- Condition: This symptom occurs if the **port-security enable** or **port-security port-mode** command is repeatedly executed.
- Remarks: None.

202207121416

- Symptom: IS-IS neighbors are disconnected during an ISSU.
- Condition: This symptom might occur if the device has established IS-IS neighbor relationships and an ISSU is performed to upgrade the software from 27xx to 67xx.
- Remarks: None.

202209200164

- Symptom: The device reports that L2VPN resources are insufficient and the AC configuration fails to be deployed.
- Condition: This symptom occurs if the following operations are performed:
 - a. When an L3VNI is created on the device, the statistics collection feature is enabled by default. Then, repeatedly create and delete L3VNIs.
 - b. Repeatedly execute the **interface vsi-interface** command and then the **undo interface vsi-interface** command.
- Remarks: None.

202209120087

- Symptom: A QoS policy that contains multiple class-behavior associations is applied to the outbound direction of the device. When the actions in a class-behavior association are modified, traffic might match another class-behavior association by mistake.
- Condition: This symptom occurs if the following operations are performed:
 - a. Apply a QoS policy to multiple interfaces. A behavior contains the counting or CAR action.
 - b. Modify the actions in a traffic behavior or match criteria in a traffic class in the QoS policy or another QoS policy. Or, apply the QoS policy again.
- Remarks: None.

202207081494

- Symptom: A process becomes abnormal when a VSI interface is assigned to a VSI.
- Condition: This symptom might occur if the VSI is associated with an EVPN instance by using the **evpn encapsulation vxlan binding instance** command.
- Remarks: None.

202109131526

- Symptom: Untagged packets cannot be forwarded for a local VLAN to a remote VXLAN.
- Condition: This symptom might occur if the device is operating in border mode and forwards untagged packets of a local VLAN over a VXLAN tunnel.
- Remarks: None.

202208220621

- Symptom: In an EVPN network, traffic loss persists between the leaf and spine tiers.
- Condition: This symptom might occur if multiple ECMP paths over VXLAN tunnels exist between the leaf tier and traffic destination, the next hops are different spine devices, and one of the spine devices is rebooted.
- Remarks: None.

202209191590

- Symptom: In an MVXLAN network, public network multicast routes created for PIM are not deleted completely.

- Condition: This symptom might occur if the MVXLAN IPv4 address family and the MVXLAN IPv6 address family are created, and then the MVXLAN IPv4 address family and the MVXLAN IPv6 address family are deleted in turn.
- Remarks: None.

202208311310

- Symptom: IPv6 automated device deployment is interrupted.
- Condition: This symptom might occur if the device performs IPv6 automated device deployment.
- Remarks: None.

202208261550

- Symptom: In an EVPN VXLAN network that forwards Layer 3 multicast traffic over DCI connections, it takes longer time than expected for traffic loss to stop after traffic flows enter the network.
- Condition: This symptom might occur if an EVPN VXLAN network forwards Layer 3 multicast traffic over DCI connections.
- Remarks: None.

202207080423

- Symptom: MAC authentication users flap on an aggregate interface 8 minutes after they come online.
- Condition: This symptom might occur if MAC authentication user offline detection is enabled by default.
- Remarks: None.

202206291177

- Symptom: The device receives NA packets that do not carry the target link-layer address field and does not learn ND entries from the NA packets.
- Condition: This symptom might occur if the device receives unrequested NA packets that do not carry the target link-layer address field.
- Remarks: None.

202206230765

- Symptom: The device reports a permission deny error.
- Condition: This symptom might occur if command authorization is enabled and the **repeat** command is executed for more than 1000 times.
- Remarks: None.

202206060838

- Symptom: In Layer 3 multicast on a cascaded M-LAG network, IGMP packets are looped between M-LAG interfaces.
- Condition: This symptom occurs if an M-LAG interface receives IGMP query packets.
- Remarks: None.

202205191660

- Symptom: In a multicast VXLAN network, the state of the MTunnel interface is down.
- Condition: This symptom might occur if you configure multicast VXLAN and then configure PIM on an interface in a VPN instance.
- Remarks: None.

202210250334

- Symptom: The number of free resources in the **display resource-monitor resource nexthoppool1** command output increases all the time, and a resource alarm is triggered
- Condition: This symptom occurs if the switch learns a large number of ARP entries and you execute the **reset arp** command.
- Remarks: None.

202209200820

- Symptom: Memory is leaked.
- Condition: This symptom occurs if you add and delete Layer 3 aggregate subinterfaces.
- Remarks: None.

202201171691

- Symptom: A QoS policy is still in effect after it is removed from a VSI interface.
- Condition: This symptom occurs if you perform the following operations:
 - a. Create a QoS policy without class-behavior associations, and apply it to a VSI interface.
 - b. Configure a class-behavior association in the QoS policy, and remove the QoS policy from the VSI interface.
- Remarks: None.

202112270288

- Symptom: In an IRF fabric with multichassis aggregation, the memory is exhausted, and the switch reboots when a large number of MAC authentication users come online on an aggregate interface.
- Condition: This symptom occurs if offline detection and reauthentication are enabled.
- Remarks: None.

202104220925

- Symptom: On an EVPN VTEP with conversational learning disabled for remote MAC address entries, a blackhole MAC address entry on an interface module is deleted even through traffic matches the entry.
- Condition: This symptom might occur if some blackhole MAC address entries are not synchronized between cards.
- Remarks: None.

Resolved problems in F6705

202206010870

- Symptom: In a network with two IRF fabrics, BFD MAD flaps.
- Condition: This symptom occurs if the following operations are performed:
 - a. Enable BFD MAD on interfaces in the same VLAN.
 - b. Perform a master/subordinate switchover on one IRF fabric.
- Remarks: None.

202204151725

- Symptom: In an MPLS L3VPN, packets are forwarded across VPN instances.
- Condition: This symptom occurs if an interface bound to VPN instance a receives a packet of VPN instance b and the packet matches the routing table of VPN instance b.

- Remarks: This issue has been solved for VLAN interfaces and Layer 3 Ethernet subinterfaces. This issue cannot be solved for Layer 3 Ethernet interfaces.

202204290654

- Symptom: In an IRF fabric with multichassis link aggregation, some of the aggregation member ports cannot forward traffic, causing uneven hashing after member ports are shut down and then brought up.
- Condition: This symptom occurs if the aggregate interface acts as an outgoing interface for a VXLAN tunnel.
- Remarks: None.

202204191568

- Symptom: The convergence time of the Monitor Link down function is long.
- Condition: This symptom occurs when the downlink interfaces in a monitor link group are shut down because an uplink interface goes down.
- Remarks: None.

202204230202

- Symptom: A MAC address move does not trigger an ND move.
- Condition: This symptom might occur in an underlay M-LAG network if the **mac-address mac-move fast-update** command is executed.
- Remarks: None.

202109131526

- Symptom: The device cannot forward untagged packets from a VLAN to a remote VXLAN.
- Condition: This symptom occurs if the device in border mode forwards untagged packets from a VLAN out of a VXLAN tunnel.
- Remarks: None.

202205091702

- Symptom: On an EVPN DRNI network, packets are dropped unexpectedly.
- Condition: This symptom occurs if packets are received on an AC, the packets are VXLAN-encapsulated, and the packets carry the PVID VLAN tag of the interface hosting the AC.
- Remarks: None.

202205091701

- Symptom: When ARP/ND traffic exists on an EVPN+DRNI network, if you repeatedly shut down and bring up the IPL aggregate interface of a DR member device, the device will reboot unexpectedly.
- Condition: This symptom occurs if you repeatedly shut down and bring up the IPL aggregate interface of a DR member device when ARP/ND traffic exists on an EVPN+DRNI network.
- Remarks: None.

202205091696

- Symptom: The reply to an HTTP request on a device carries the server:HTTPD field, which is used to identify the server information. The vulnerability scanners consider that the server field might disclose the server information and result in attacks.
- Condition: This symptom occurs if the device receives HTTP requests.
- Remarks: None.

202205091688

- Symptom: The memory leaks for the routed module.
- Condition: This symptom occurs if you configure a gRPC sensor path to collect route information, and then make routes on the device flap.
- Remarks: None.

202203141354

- Symptom: After the device is rebooted, the detection interval configured for the BFD echo session does not take effect, and is displayed as the default value.
- Condition: This symptom occurs if the following operations are performed on a DRNI network:
 - a. Configure a static BFD echo session with a detection interval different from that configured for the BFD echo session on an interface. The session can be negotiated as up.
 - b. Save the configuration, and then reboot the device.
- Remarks: None.

202205171718

- Symptom: When identical static ARP entries are configured on the DR member devices in a DR system, configuration fails on one DR member device.
- Condition: This symptom might occur if identical static ARP entries are configured on the DR member devices in a DR system.
- Remarks: None.

202105150186

- Symptom: After an aggregate interface authenticates a MAC authentication user, an IRF master/subordinate switchover occurs, and the user goes offline 10 minutes later.
- Condition: This symptom occurs if an aggregate interface authenticates a MAC authentication user and an IRF master/subordinate switchover occurs.
- Remarks: None.

202206240523

- Symptom: A downstream device attached to an M-LAG EVPN system cannot ping the gateway.
- Condition: This symptom occurs if the VSI interfaces acting as gateways are assigned different IP addresses and the downstream device pings one gateway address.
- Remarks: None.

202204290651

- Symptom: Layer 3 aggregate subinterfaces do not forward traffic.
- Condition: This symptom might occur if cross-device aggregation is configured in stack deployment and both Layer 3 aggregate subinterfaces and Layer 3 subinterfaces act as equal-cost outgoing interfaces for a VXLAN tunnel.
- Remarks: Shut down and bring up any outgoing interface for the VXLAN tunnel after patch installation.
- Remarks: None.

202204110097

- Symptom: In an MPLS VPLS network, the PE device cannot transparently transmit IGMP packets when multicast features are enabled on the PE device.
- Condition: This symptom occurs when the device acts as a PE of an MPLS VPLS network, and Layer 2 or Layer 3 multicast features are enabled on the device.
- Remarks: None.

202103051212

- Symptom: On a DR system, EVPN VXLAN fails to forward untagged traffic from the local site to a remote site.
- Condition: This symptom might occur if the untagged packets are incorrectly processed when they are forwarded at Layer 3 out of a tunnel interface on the underlay network.
- Remarks: None.

Resolved problems in E6702

None.

Resolved problems in R2717

202107141294

- Symptom: A DR interface does not learn MAC addresses.
- Condition: This symptom might occur if the following conditions exist:
 - The IPP is an aggregate interface.
 - In system view, the **mac-address mac-learning enable** command is executed, undone, and then executed again.
- Remarks: None.

202107120224

- Symptom: A DR member device learns some MAC addresses on incorrect ports.
- Condition: This symptom might occur if the MAC addresses move between the DR member devices.
- Remarks: None.

202107050018

- Symptom: On an IRF fabric, downstream aggregate interfaces fail, reducing ECMP links. As a result, it takes more time to finish route convergence.
- Condition: This symptom occurs if you manually shut down one or multiple downstream aggregate interfaces on the IRF fabric.
- Remarks: None.

202105260239

- Symptom: The device is operating in IPv6_128 mode and reports insufficient ND resources when the host entry resources are exhausted.
- Condition: This symptom occurs when the ARP table capacity and the 32-bit routing table capacity exceed 16K.
- Remarks: None

202105201209

- Symptom: NS requests are flooded by EVPN in a Layer 2 network if their target IP address belongs to a distributed EVPN gateway and destination IP address is a multicast address.
- Condition: This symptom might occur if local ND proxy is enabled on the distributed EVPN gateway.
- Remarks: None

202105200314

- Symptom: The device reports an incorrect message of insufficient PBR resources.
- Condition: This symptom occurs if you specify an overlay ECMP next hop for the service chain PBR policy on the device.
- Remarks: None.

202105171205

- Symptom: In NLB scenario, a member cannot ping another member.
- Condition: This symptom occurs when ARP suppression is enabled on the device to directly answer ARP requests and the attached VMs form a cluster. The MAC address in the Ethernet header of the ARP request is the cluster MAC address. The protocol's MAC address is the real MAC address of the sender. The MAC address in the Ethernet header of the ARP reply is also the cluster MAC address. Therefore, traffic cannot be forwarded.
- Remarks: None.

202105171186

- Symptom: Traffic cannot be forwarded after ARP entries moved.
- Condition: This symptom occurs in an EVPN DRNI network where the DRNI device's AC-side and tunnel-side ARP entries moved.
- Remarks: None.

202105170291

- Symptom: uRPF cannot be enabled or disabled on an aggregate interface when the device is operating in I2gw hardware resource mode.
- Condition: None.
- Remarks: Enable or disable uRPF again after the patch is installed.

202104291407

- Symptom: Packets are dropped on 100-GE IRF physical interfaces because of CRC errors.
- Condition: This symptom might occur if the 100-GE IRF physical interfaces are connected through DAC cables.
- Remarks: None.

202106160130

- Symptom: The resmond process gets stuck. As a result, no output is displayed when you execute the **display resource-monitor** command.
- Condition: This symptom might occur if the controller frequently deploys resource management subscriptions related to the resmond process and cancels the subscriptions.
- Remarks: None.

202108240145

- Symptom: The function of allowing a Layer 2 aggregate interface to be the only member of the specified DR group will cause packet loss.
- Condition: This symptom occurs if the controller deploys the configuration of assigning an aggregate interface to a DR group and the controller enables the function of allowing a Layer 2 aggregate interface to be the only member of the specified DR group by default.
- Remarks: None.

202101141489

- Symptom: The ACL specifications for MAC authentication cannot meet the customer requirements.

- Condition: This symptom occurs if MAC authentication deploys a large number of ACLs.
- Remarks: Optimize ACL resource deployment to save resources. You can save resources for only the first ACL referenced by MAC authentication on each port.

202105150409

- Symptom: It takes a long time to remove member ports from an aggregation group.
- Condition: This symptom occurs if you remove member ports from an aggregation group.
- Remarks: None.

202101060459

- Symptom: The tunnel encapsulation entry errors might occur with a low probability, which will cause traffic forwarding failure.
- Condition: This symptom occurs if you change the source interface of a tunnel when the tunnel is up.
- Remarks: None.

202008210114

- Symptom: Configure an attack defense policy, but do not apply it to the local device. When the **reset attack-defense statistics local** command is executed to clear attack detection and prevention statistics for the device, the device reboots unexpectedly.
- Condition: This symptom occurs if you configure an attack defense policy, but do not apply it to the local device.
- Remarks: None.

202108251354

- Symptom: A vulnerability is found by mistake when Nessus is used to scan vulnerabilities on the device.
- Condition: This symptom occurs if the following conditions exist:
 - Telnet is enabled on the HPE device.
 - Nessus is used to scan the Telnet service port of the device.
- Remarks: None.

202102250815

- Symptom: When an ACL is applied to an aggregate interface for packet filtering, some of its ACL rules do not take effect.
- Condition: This symptom occurs if the following operations are performed:
 - a. Apply an ACL that contains a large number of ACL rules to an aggregate interface. The ACL occupies the resources in both slice 6 and slice 7.
 - b. Apply another ACL that occupies some resources in slice 4.
 - c. Release some ACL resources in slice 6 through deleting some ACL rules.
 - d. Apply one more ACL that occupies some resources in slice 6 released in the previous step after occupying all available resources in slice 4.
- Remarks: Use the **display qos-acl resource advanced-mode** command to view the ACL resource usage.

202104120090

- Symptom: The IPP learns the MAC address of a DR interface. As a result, the ARP entry that should be learned on the DR interface is learned on the IPP.
- Condition: This symptom occurs if MAC address learning is disabled on the IPP.
- Remarks: None.

202012220599

- Symptom: The state of a port becomes down.
- Condition: This symptom occurs if the local port goes down because the link is instable or the peer port goes down.
- Remarks: None.

202103151648

- Symptom: The device redistributes traffic across the Selected ports of an aggregate interface.
- Condition: This symptom might occur if resilient load sharing is enabled on the aggregate interface and a member port of the aggregate interface is shut down.
- Remarks: None.

202103051142

- Symptom: On an IRF fabric, if you try to switch an Ethernet interface with empty configuration on an FC card in the master IRF member device to an FC interface, the system prompts that the operation failed.
- Condition: This symptom occurs if you try to switch an Ethernet interface with empty configuration on an FC card in the master IRF member device to an FC interface.
- Remarks: None.

202101210277

- Symptom: The device reboots unexpectedly.
- Condition: This symptom occurs if the **debug netstream show ipv6_agg slot 1** command is executed in probe view.
- Remarks: None.

202103011191

- Symptom: The device does not remove the router MAC attribute when reoriginating EVPN routes based on the routes received from CEs.
- Condition: This symptom might occur if the CEs advertise unicast routes that carry the router MAC attribute.
- Remarks: None.

202012160051

- Symptom: The ACL issued for MAC authentication cannot prevent packets with an unknown source MAC address from being sent to the CPU. The device does not forward those packets.
- Condition: This symptom might occur if the ACL contains deny rules.
- Remarks: None.

202103250502

- Symptom: VXLAN mapping does not take effect on a Layer 2 DCI network.
- Condition: This symptom might occur if the following tasks are performed:
 - a. Execute the **dc enable** command.
 - b. Create VXLAN-DCI tunnels.
 - c. Execute the **mapping vni** command.
- Remarks: None.

202101230067

- Symptom: Some MAC address entries do not age out on a DR system with singlehomed devices attached.

- Condition: This symptom might occur if the DR member devices receive packets sourced from the same MAC address.
- Remarks: None.

202012140818

- Symptom: Traffic forwarding fails after the tunnel source or destination IP address is modified for a large number of VXLAN tunnels.
- Condition: This symptom might occur if the tunnel source or destination IP address is modified for a large number of VXLAN tunnels.
- Remarks: None.

202103051094

- Symptom: An aggregation device is named incorrectly after it is automatically deployed.
- Condition: This symptom might occur if an aggregation device starts with initial configuration and is automatically deployed from the controller.
- Remarks: None.

202103051092

- Symptom: After the device finishes automatic IRF setup, the IRF physical interfaces and BFD MAD interfaces are not configured as expected. The candidate IRF physical interfaces become the BFD MAD interfaces, and the candidate BFD MAD interfaces become the IRF physical interfaces.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure IRF and BFD MAD settings after the device finishes automatic onboarding.
 - b. Delete the .mdb startup configuration files and reboot the device with a .txt startup configuration file.
- Remarks: None.

202103051104

- Symptom: On an aggregate interface, the ACL rules used for packet filtering do not take effect.
- Condition: This symptom might occur if the following tasks are performed:
 - a. Issue an ACL to an aggregate interface to occupy the resources of slice 6 and slice 7.
 - b. Issue an ACL to occupy some resources of slice 4.
 - c. Delete ACL rules to release some resources of slice 6.
 - d. Issue an ACL to occupy the remaining resources of slice 4 and the resources of slice 6 that have been released in step c.
- Remarks: Execute the **display qos-acl resource advanced-mode** command to view ACL resource usage.

202103051098

- Symptom: Services provided by a VM are unavailable.
- Condition: This symptom might occur if the VM migrates several times within seconds in an EVPN network.
- Remarks: None.

202103051097

- Symptom: Traffic forwarding fails after the tunnel source or destination IP address is modified for a large number of VXLAN tunnels.
- Condition: This symptom might occur if the tunnel source or destination IP address is modified for a large number of VXLAN tunnels.

- Remarks: None.

202103190424

- Symptom: The policy-based routing (PBR) configuration on a GRE tunnel is lost after the device reboots.
- Condition: This symptom occurs after you configure PBR on the GRE tunnel interface and then reboot the device.
- Remarks: None.

202103190423

- Symptom: A BFD session flap occurs.
- Condition: This symptom occurs if you execute the **ptp mode** command to specify a clock node type for the device after the BFD session comes up.
- Remarks: None.

202009181318

- Symptom: After an ARP entry ages out, the packets with the IP address in the ARP entry as the destination IP address can still be forwarded.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure a route, and specify the outgoing interface and next hop.
 - b. Send packets to the destination IP address of the static route to make the device learn the ARP entry corresponding to the next hop address.
 - c. View the routing and forwarding information after the ARP entry ages out.
- Remarks: None.

202007011330

- Symptom: A leaf node cannot forward the traffic received from ACs over a VXLAN tunnel.
- Condition: This symptom might occur if the following conditions exist:
 - The leaf node has multiple ECMP routes to the tunnel destination address, and the ECMP routes are shared by a large number of VXLAN tunnels.
 - The remote site is multihomed to leaf nodes, and the site-facing interfaces flap on the leaf nodes.
- Remarks: None.

202008210714

- Symptom: The **tracert** command does not output the complete route across an EVPN network.
- Condition: This symptom might occur if the **tracert** command is executed on a leaf node whose attached devices access the external network through the border leaf node.
- Remarks: None.

202007080188

- Symptom: After a DR system is set up, a DR interface stays in STP discarding state, which causes forwarding failure.
- Condition: This symptom might occur if the spanning tree feature is enabled on the primary member in the DR system.
- Remarks: None.

202009110593

- Symptom: VSI bandwidth restraints do not take effect.
- Condition: This symptom might occur if the following operations are performed:

- a. Execute the **restrain broadcast 0**, **restrain multicast 0**, or **restrain unknown-unicast 0** command on a VSI.
 - b. Reboot the switch with the configuration.
- Remarks: None.

202008111693

- Symptom: After a primary/backup LSP switchover, the numbers of primary and backup LSPs become inconsistent with those before the switchover.
- Condition: This symptom might occur after a primary/backup LSP switchover in the stacking network environment.
- Remarks: None.

202006080578

- Symptom: When the switch is connected to a Juniper router by using a BFD-enabled aggregate link, the aggregate link cannot recover after interface flapping occurs.
- Condition: This symptom might occur if the switch is connected to a Juniper router by using a BFD-enabled aggregate link and the BFD session goes down.
- Remarks: None.

202007030035

- Symptom: A short physical state change suppression interval causes undesired interface flapping.
- Condition: This symptom might occur if the switch is installed with multiple modules.
- Remarks: None.

202006221179

- Symptom: Incoming packet statistics for a VSI interface are incorrect.
- Condition: This symptom might occur if packet statistics collection is enabled and then disabled for a VSI interface.
- Remarks: None.

202009100495

- Symptom: The switch might fail to respond to **display** commands.
- Condition: This symptom might occur if the following conditions exist:
 - Accounting is enabled, but the accounting server is unreachable.
 - Some **display** commands are executed repeatedly.
- Remarks: None.

202006100478

- Symptom: A VXLAN VTEP cannot ping a directly connected port on a peer.
- Condition: This symptom might occur if the **hardware-resource vxlan l2gw** command is executed and QinQ is enabled on the interface connected to the peer.
- Remarks: None.

202009010524

- Symptom: In a VXLAN network, the **display l2vpn mac-address** command cyclically displays the first 24 MAC address entries.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure more than 4000 VSIs.
 - b. Configure a VSI with an index larger than 4000 to learn more than 24 MAC addresses.

- c. Execute the **display l2vpn mac-address** command.
- Remarks: You can press Ctrl+C to exit the cyclical display.

202005150380

- Symptom: The controller cannot discover a VTEP or set up VXLAN tunnels for the VTEP.
- Condition: This symptom might occur if the VTEP boots with the **tunnel global source-address** setting.
- Remarks: None.

202009240007

- Symptom: When the device issues IPv6 addresses to a VXLAN IP gateway, the gateway displays link-local addresses of VSI interfaces as **duplicate** and global unicast addresses as **tentative**.
- Condition: This symptom might occur if the **distributed-gateway local** command is executed on VSI interfaces and then the **vxlan tunnel nd-learning disable** command is executed in system view.
- Remarks: None.

202005210724

- Symptom: On an EVPN network, the routes reflected by a route reflector do not carry L3 VXLAN IDs.
- Condition: This symptom might occur if edge devices on the EVPN network frequently generate route updates.
- Remarks: None.

202004081642

- Symptom: The memory usage is high.
- Condition: This symptom might occur if the following operations are performed:
 - Telnet to the switch. Execute the **ftp** command without specifying an IP address.
 - Close the Telnet window.
 - Repeat steps a and b.
- Remarks: None.

202010100330

- Symptom: When both EVPN and DRNI are configured, underlying resource **nexthop pool 1** is not deleted completely after certain operations are performed.
- Condition: This symptom might occur if the following operations are performed in a DRNI network:
 - Create a DR interface and configure ACs on it.
 - Delete the DR interface.
- Remarks: None.

202009180971

- Symptom: The traffic accounting units change from bytes and packets to only bytes in the traffic accounting action of an applied QoS policy.
- Condition: This symptom occurs if the following operations are performed:
 - Configure two QoS policies. In the two QoS policies, associate different traffic classes with the same traffic behavior that is configured with the traffic accounting action. The traffic accounting action counts packets in both bytes and packets.

- b. Apply one QoS policy in the inbound direction of an interface, and apply the other QoS policy in the outbound direction of the interface.
 - c. The traffic accounting units of the first applied QoS policy change from bytes and packets to only bytes.
- Remarks: None.

202006240947

- Symptom: When you apply a QoS policy, the system prompts that the QoS and ACL resources are insufficient.
- Condition: This symptom occurs if the traffic classifiers of the QoS policy reference both IPv4 and IPv6 ACLs.
- Remarks: None.

202009070887

- Symptom: The OSPF packets received on ACs might be dropped with a low probability.
- Condition: This symptom might occur if the following operations are performed:
 - a. Connect an IRF fabric to a peer IRF fabric through ECMP links.
 - b. Establish tunnels A and B between the two IRF fabrics. The next hops of tunnels A and B are on different IRF member devices.
 - c. The instant aging of the ARP entry for tunnel A or B triggers ARP updates.
- Remarks: None.

202007230547

- Symptom: When both EVPN and DRNI are configured, the switch advertises MAC/IP advertisement routes that carry incorrect nexthops.
- Condition: This symptom might occur if a server is attached to a DR interface of the switch and the MAC/IP advertisement routes carry information about the server.
- Remarks: None.

202009300174

- Symptom: DGPP entries are not synchronized among member devices in an IRF fabric after the IRF fabric reboots.
- Condition: This symptom occurs if the following conditions exist:
 - A multichassis Ethernet aggregate interface is configured as the next hop output interface of a tunnel.
 - The IRF fabric reboots with the configuration.
- Remarks: None.

202004081619

- Symptom: The device cannot be logged in.
- Condition: This symptom occurs if password control is enabled on the device and the system time change causes the login password to expire.
- Remarks: None.

202008210765

- Symptom: On an IRF fabric deployed with PBR policies, resources on a subordinate device are insufficient and packet forwarding exceptions occur on the device.
- Condition: This symptom might occur if the following conditions exist:
 - Service chain rules are configured in PBR actions.
 - The next hops of PBR policies are flapping.

- Remarks: If this issue has occurred on a subordinate device before the patch is installed, you must reboot the subordinate device to resolve the issue from the device.

202007080258

- Symptom: On a DR system that uses an Ethernet aggregate link as the IPL, a DR member device cannot forward the traffic received on a DR interface out of the IPP.
- Condition: This symptom might occur if member ports of the IPP and the DR interface are in Unselected state, and the DR interface is deleted by using the **undo interface Bridge-Aggregation** command.
- Remarks: none.

202008140876

- Symptom: The time in the **display clock** command output is not accurate.
- Condition: This symptom occurs if the following conditions exist:
 - The **clock protocol ntp** command is executed to specify NTP for obtaining the time.
 - The time difference between the system and the NTP server exceeds 68 years.
- Remarks: Install the patch and then cold reboot the device, or configure the system time to reduce its difference with the NTP server's time to less than 68 years by executing the following commands in order:
 - clock protocol none**
 - clock datetime 00:00:00 2000/11/11**
 - clock protocol ntp**
- Remarks: none.

202006280209

- Symptom: The number of received packets and the number of sent packets on an interface abnormally increase in the interface statistics.
- Condition: This symptom occurs if a 40-Gbps transceiver module is removed from a 100-GE interface.
- Remarks: none.

202007070356

- Symptom: The device might reboot unexpectedly.
- Condition: None.
- Remarks: none.

202005070894

- Symptom: The switch cannot ping a server that is connected to an interface AC.
- Condition: This symptom might occur if an interface is removed from an aggregation group and then configured as an AC.
- Remarks: none.

202009220586

- Symptom: On a distributed EVPN gateway, traffic is not forwarded to the expected next hop based on a PBR policy.
- Condition: This symptom might occur if VSI interfaces receive continuous ARP request traffic.
- Remarks: none.

202006181187

- Symptom: Lower-layer ACL resources are not released.
- Condition: This symptom occurs if the following operations have been performed:
 - Apply an IPv6 ACL without rule 0 to an Ethernet interface.

b. Remove the IPv6 ACL from the interface.

- Remarks: If the issue has occurred on a device, installing this patch cannot automatically resolve the issue, and you must remove the ACL from the interface and then reapply it.

202010150963

- Symptom: The **reset packet-drop** command cannot clear the dropped packet statistics for an interface.
- Condition: This symptom occurs if the **reset packet-drop** command is executed to clear the dropped packet statistics when congestion occurs on an interface.
- Remarks: None.

202005220490

- Symptom: In a VXLAN network, traffic from VMs cannot be forwarded.
- Condition: This symptom occurs if an IRF fabric acts as a VTEP, the IRF fabric is configured with an AC, and then the subordinate IRF member device is rebooted.

202002251001

- Symptom: No error message is prompted for patch installation failure.
- Condition: This symptom occurs if you log in to the device through Telnet or SSH and the patch installation fails.

202005191016

- Symptom: A 10-GE transceiver module inserted into a 40-GE interface by using a 40-GE to 10-GE adapter fails to transmit optical signals correctly.
- Condition: This symptom occurs if a 10-GE transceiver module is inserted into a 40-GE interface by using a 40-GE to 10-GE adapter.
- Remarks: None.

202005210484

- Symptom: After the operating mode of a copper port on an LSWM124XGT2Q card is switched between Layer 2 and Layer 3, the port cannot come up.
- Condition: This symptom occurs if the copper port is configured to operate at 1 Gbps and the operating mode of the port is repeatedly switched between Layer 2 and Layer 3.
- Remarks: None.

202005090333

- Symptom: After you configure a PBR policy and enable packet statistics for a Layer 3 Ethernet subinterface, the PBR policy cannot take effect.
- Condition: This symptom might occur if you configure a PBR policy and enable packet statistics for a Layer 3 Ethernet subinterface.
- Remarks: None.

202004090704

- Symptom: CRC errors might occur on a 100-GE transceiver module, causing packet loss.
- Condition: This symptom might occur if a 100-GE interface uses a 100-GE transceiver module for connection.
- Remarks: None.

202004300168

- Symptom: For a 40-GE interface manually shut down, a 10-GE transceiver module inserted into this interface by using a 40-GE to 10-GE adapter can transmit optical signal correctly. After the transceiver module is removed and reinstalled in the 40-GE interface, the interface comes up.

- Condition: This symptom occurs when the following operations have been performed:
 - a. Execute the **shutdown** command on the 40-GE interface.
 - b. Insert a 40-GE to 10-GE adapter into the 40-GE interface.
 - c. Insert a 10-GE transceiver module into the adapter and connect the interface to a peer device.
 - d. Remove and reinstall the 10-GE transceiver module in the interface.
- Remarks: None.

202004071075

- Symptom: The remote copy link between the switch and a 3PAR device flaps constantly.
- Condition: This symptom might occur if the following conditions exist:
 - The switch uses ports on the LSWM116FC interface card to connect to 3PAR devices running OS 3.3.1 MU2.
 - 3PAR peer persistence is configured.
 - A 3PAR device has a power outage and then recovers.
- Remarks:
 - This software version can resolve the link flapping issue when 3PAR devices run OS 3.3.1 MU2. It is not verified whether this issue can be resolved when 3PAR devices run any other OS versions.
 - If the device has been upgraded to F2707, you must execute the **undo fc name-service auto-discovery** command on the VSAN that accommodates the interfaces with remote copy links attached.

202004231154/202004240282

- Symptom: In a VRRP group, the device with higher priority is elected as the backup and cannot become the master.
- Condition: This symptom might occur if you continuously modify the device priorities to perform master/backup switchover in the VRRP group (with version VRRPv2 or VRRPv3).
- Remarks: None.

202004081694

- Symptom: The memory usage is too high.
- Condition: This symptom occurs after the following operations have been performed:
 - a. Use the **snmp-agent trap enable syslog** command to enable sending SNMP traps for system log messages.
 - b. Use the **info-center syslog trap buffersize** command to set the maximum number of log traps that can be stored in the log trap buffer.
- Remarks: None.

202004241070

- Symptom: On a multicast network, the interval between two IGMP reports received by the IGMP querier is 0.5 seconds longer than the maximum response time for IGMP general queries set on the device.
- Condition: This symptom occurs if the receiver-side device is enabled with IGMP snooping proxying.
- Remarks: None.

202004230866

- Symptom: Aggregation group membership change might cause VXLAN traffic loss.

- Condition: This symptom might occur if the aggregate interface of the aggregation group that has membership change is the outgoing interface for a VXLAN tunnel.

202004290297

- Symptom: The match order of issued PBR policy nodes is incorrect.
- Condition: This symptom might occur if PBR policies are issued to multiple interfaces and the interface (pointing to a next hop in a PBR policy) in an ARP entry has change to another interface.
- Remarks: None.

202005060911

- Symptom: On a data center network, the device (a leaf node) fails to learn the ARP entry for a failed VM upon receiving an ARP request from the VM, which cause traffic forwarding failure.
- Condition: This symptom occurs if the following conditions exist:
 - The device is configured with a large number of ACs and has a large number of security policy-related OpenFlow flow entries.
 - OpenFlow connections repeatedly flap when the device learns a large number of ARP entries.
- Remarks: None.

202004290738

- Symptom: IPv4 or IPv6 Layer 3 VPN traffic is interrupted when the public network routes repeatedly flap on an IRF fabric.
- Condition: This symptom might occur if the following conditions exist:
 - On the IRF fabric, a multichassis aggregate interface acts as the output interface of BGP public network routes.
 - The member ports of the aggregate interface are repeatedly shut down and then brought up.
- Remarks: None.

202004301450

- Symptom: A patch or feature image cannot be uninstalled from an IRF subordinate device.
- Condition: This symptom occurs if the installation path of the patch or feature image on the subordinate device is inconsistent with that on the master device.
- Remarks: None.

202003190521

- Symptom: A traffic classifier that references an ACL to match the inner header information of VXLAN packets does not classify traffic based on the ACL.
- Condition: This symptom might occur if a traffic classifier is configured to reference an ACL to match the inner header information of VXLAN packets by using the **if-match acl name inner** command.
- Remarks: None.

202004290414

- Symptom: The return packets are not forwarded over the same path as the originating packets.
- Condition: This symptom occurs if the originating packets are MPLS packets and are decapsulated before being forwarded at Layer 3.
- Remarks: None.

202005070894

- Symptom: In a VXLAN network, VMs in the same subnet might fail to ping one another after certain operations are performed.
- Condition: This symptom might occur if the following operations are performed:
 - a. Create an Ethernet service instance on an interface, configure the SVLAN ID match criterion for it, and set the PVID of that interface to be the same as the SVLAN ID in the match criterion.
 - b. Delete the Ethernet service instance.
 - c. Repeat steps a and b.
- Remarks: None.

202003181317/201912171153

- Symptom: Few MPLS packets are dropped.
- Condition: This symptom occurs if the routes are updated after session protection is enabled in LDP view.
- Remarks: None.

202001130806

- Symptom: When executing the **irf member renumber** command, the system should output a message indicating that a reboot is required for this command to take effect. However, the system does not output this message.
- Condition: This symptom occurs when the **irf member renumber** command is executed.
- Remarks: None.

201907300860

- Symptom: The device abnormally reboots because of PCIE access exception.
- Condition: This symptom might occur when a PCIE access exception occurs.
- Remarks: None.

201912260195

- Symptom: 10-GE ports on the local device are connected to the breakout interfaces of a 40-GE port on the neighbor device through AOC cables. Packet loss occurs on all the 10-GE ports connected to the breakout interfaces.
- Condition: This symptom occurs if you remove and then insert the AOC cable for one of the 10-GE breakout interfaces on the neighbor device.
- Remarks: None.

202001130415/201911210226

- Symptom: Error exists on the runts statistics (the number of super small frames) for IRF physical interfaces.
- This symptom occurs if the following conditions exist:
 - The IRF fabric contains two member devices.
 - The IRF fabric has forwarded frames across IRF member devices.
 - The IRF fabric counts statistics for IRF physical interfaces.
- Remarks: None.

202002060416

- Symptom: BFD MAD still remains in Faulty state on an IRF fabric after the IRF fabric recovers from an IRF split event.
- Condition: This symptom occurs if the following conditions exist:

- a. The IRF fabric contains two member devices and BFD MAD is configured on the IRF fabric.
- b. The IRF fabric splits and then recovers.
- Remarks: None.

202001100390

- Symptom: The command that splits an interface into four breakout interfaces is stuck.
- Condition: This symptom might occur if the commands that split and combine interfaces are executed repeatedly in a VXLAN network.
- Remarks: None.

202001070591

- Symptom: The destination MAC address is incorrect for packets sent out of a transport-facing interface on a leaf device in a VXLAN network.
- Condition: This symptom might occur if the transport-facing interface of the leaf device has multiple next hops and the next hops have changes.
- Remarks: None.

201905210862

- Symptom: The authentication server is stuck in **block** state and cannot restore to **active** state.
- Condition: This symptom might occur if two master/backup device switchovers occur before the authentication server placed in **block** state recovers to **active** state.
- Remarks: None.

202001190271

- Symptom: The telnet operation hangs with a low probability.
- Condition: This symptom might occur if you telnet to the device, and enable command accounting but the accounting server is not available.
- Remarks: None.

201905210848

- Symptom: The link aggregation module cannot process services when the BFD session flaps.
- Condition: This symptom might occur if you configure collaboration between Ethernet link aggregation and BFD.
- Remarks: None.

202002180298

- Symptom: The packet statistics for VLAN interfaces and VSI interfaces are incorrect.
- Condition: This symptom occurs if packet statistics are collected for VLAN interfaces and VSI interfaces.
- Remarks: None.

202001140870

- Symptom: The switch reboots because of memory exhaustion.
- Condition: This symptom might occur if the memory management unit has parity errors.
- Remarks: If infinite loops cannot be recovered, do not upgrade the switch to this software version.

202001060626

- Symptom: On an OVSDB network, the controller might mistakenly delete the Ethernet service instance configuration of interfaces.

- Condition: This symptom occurs if the following operations have been performed:
 - a. Use the Nuage controller on the network. The devices form an IRF fabric to act as a VTEP.
 - b. Execute the **vtep access port** command on interfaces of the active MPU and standby MPU.
 - c. Perform a master/subordinate switchover on the IRF fabric.
- Remarks: None.

201912181125

- Symptom: Few MPLS packets are dropped.
- Condition: This symptom occurs if the routes are updated after session protection is enabled in LDP view.
- Remarks: None.

201912300910

- Symptom: When the automatic configuration feature is used to replace an IRF member device, the IRF member devices not replaced also reboot during the replacement process.
- Condition: This symptom occurs when the automatic configuration feature is used to replace an IRF member device.
- Remarks: None.

202003030727

- Symptom: The VMs not in the CIDR range can access the external network.
- Condition: This symptom might occur if the following operations have been performed:
 - a. Configure a CIDR network on Device A. The CIDR network does not contain the internal network of Device A.
 - b. Configure the same RTs for the BGP IPv4, IPv6, and EVPN address families on Device A.
 - c. The hosts on the internal network of Device A advertise routes containing the above-mentioned RTs.
- Remarks: None.

201911270018

- Symptom: gRPC fails to push data of the lfmgr/Statistics table on the device to collectors in seconds.
- Condition: This symptom occurs if gRPC pushes data of the lfmgr/Statistics table on the device to collectors.
- Remarks: None.

201911110471

- Symptom: The CLI is stuck if the **display interface** command is executed after the device is rebooted.
- Condition: This symptom occurs if the **display interface** command is executed after the device is rebooted.
- Remarks: None.

201908060060

- Symptom: The help information for the **display interface** command cannot be displayed.
- Condition: This symptom occurs if the **ifmgr** process is restarted.
- Remarks: None.

201908220593

- Symptom: The libssh2 function has an input validation vulnerability.
- Condition: This symptom occurs if the device executes remote commands or transfers files.
- Remarks: None.

201912201076

- Symptom: The device reboots unexpectedly when the **display interface fc x/x/x** command is executed to display FC interface information.
- Condition: This symptom occurs if an Ethernet interface on the device is switched to an FC interface and then the **display interface fc x/x/x** command is executed to view FC interface information.
- Remarks: None.

201912170482

- Symptom: After a reboot, the switch cannot forward VXLAN traffic based on a static route, and a static ARP entry becomes a blackhole entry.
- Condition: This symptom might occur if the following operations are performed on the switch:
 - a. Configure a static ARP entry.
 - b. Save the running configuration.
 - c. Reboot the switch.
- Remarks: None.

201912130094

- Symptom: The switch operates incorrectly after being rebooted with certain DRNI settings.
- Condition: This symptom might occur if the following operations are performed on the switch:
 - a. Associate a VPN instance with a VXLAN tunnel used by DRNI.
 - b. Save the running configuration.
 - c. Reboot the switch.
- Remarks: None.

201912120037

- Symptom: An IRF fabric cannot forward TRILL-encapsulated unknown multicast packets or forwards duplicate TRILL-encapsulated unknown multicast packets.
- Condition: This symptom might occur if the TRILL-enabled IRF fabric sends TRILL-encapsulated unknown multicast packets to a downstream device connected to the IRF fabric by using a multichassis link aggregation
- Remarks: None.

201910300793

- Symptom: The **vtep access port** command configuration is lost on an aggregate interface.
- Condition: This symptom occurs after you modify the system name of the server after the switch comes online through automatic deployment.
- Remarks: None.

201910300835

- Symptom: The **apply as-path** command cannot be issued through a NETCONF merge operation if this setting already exists on the switch.
- Condition: This symptom might occur if the **apply as-path** command is issued through a NETCONF merge operation and the switch already has this setting.

- Remarks: None.

201907181172

- Symptom: On an EVPN VTEP that is a member of a DR system, a MAC address flaps between an access switch-facing DR interface and a transport-facing interface.
- Condition: This symptom might occur if the following conditions exist on the VTEP:
 - The IPL is a VXLAN tunnel.
 - An access switch is attached to the DR interface.
 - The access switch receives ARP packets from the site attached to it.
- Remarks: None.

201908300161

- Symptom: An HPE FlexFabric 5940 2-slot Switch (JH397A) or HPE FlexFabric 5940 4-slot Switch (JH398A) switch reboots unexpectedly.
- Condition: This symptom occurs if the following operations have been performed.
 - a. Configure an LSWM124XX interface module on the switch without installation of the interface module.
 - b. Install an LSWM124XX interface module on the switch after the pre-configuration is complete.
- Remarks: None.

201910231197

- Symptom: Users are logged off from the vBRAS IRF fabric connected to the device.
- Condition: This symptom occurs if the following conditions exist:
 - The vBRAS IRF fabric is connected to the device through an aggregate interface.
 - LACP MAD is enabled on the aggregate interface by using the **mad enable** command on the vBRAS IRF fabric.
 - The vBRAS IRF fabric splits.
- Remarks: None.

201911040571

- Symptom: Failed to create a VSI interface by using the **interface vsi** command.
- Condition: This symptom might occur if the following operations are performed:
 - a. Create a service loopback group and assign member ports to the service loopback group.
 - b. Create GRE tunnel interfaces.
 - c. Create a VSI interface.
- Remarks: None.

201910080448

- Symptom: Transient packet loss occurs on an interface when the **undo packet-filter** command is executed to remove an ACL from the interface.
- Condition: This symptom might occur if the ACL has multiple rules and the action is set to deny in the last rule.
- Remarks: None.

201909170277

- Symptom: On an OVSDB network, the controller might mistakenly delete the Ethernet service instance configuration of interfaces.
- Condition: This symptom occurs if the following operations have been performed:

- a. Use the Nuage controller on the network. The devices form an IRF fabric to act as a VTEP.
- b. Execute the vtep access port command on interfaces of the active MPU and standby MPU.
- c. Perform a master/subordinate switchover on the IRF fabric
- Remarks: None.

201909290374

- Symptom: The connection between the controller and the device is disconnected.
- Condition: This symptom might occur if the controller receives a large number of packets.
- Remarks: None.

201907290489

- Symptom: The host cannot ping the gateway that has a PBR policy configured.
- Condition: This symptom might occur when you ping the switch (acting as the gateway) configured with a PBR policy from the host.
- Remarks: None.

201907181172

- Symptom: On an EVPN VTEP that is a member of a DR system, a MAC address flaps between an access switch-facing DR interface and a transport-facing interface.
- Condition: This symptom might occur if the following conditions exist on the VTEP:
 - The IPL is a VXLAN tunnel.
 - An access switch is attached to the DR interface.
 - The access switch receives ARP packets from the site attached to it.
- Remarks: None.

201907110050

- Symptom: Service traffic on an interface of the switch cannot be correctly forwarded.
- Condition: This symptom might occur if the following conditions exist:
 - The switch is connected to a VCF controller.
 - The interface is deleted or the card where the interface resides was uninstalled.
- Remarks: None.

201906060558

- Symptom: An interface configured with a PBR policy flaps and the PBR policy no longer takes effect when ECMP is configured on the interface.
- Condition: This symptom might occur if ECMP is configured on an interface where a PBR policy is applied.
- Remarks: None.

201908260252

- Symptom: In an EVPN network with distributed gateways configured, when multiple HPE 5940 switches join the network as leaf nodes, the EVPN service on the existing devices might be affected.
- Condition: This symptom might occur if multiple HPE 5940 switches are configured as leaf nodes of an EVPN network where distributed gateways are deployed.
- Remarks: None.

201905141113/201901070710

- Symptom: Some tunneled packets are lost on the output interface.

- Condition: This symptom occurs when the output interface for tunneled packets changes from a physical interface to an aggregate interface.
- Remarks: None.

201907030138

- Symptom: The types of some interfaces are not displayed in the **display interface ?** command output.
- Condition: This symptom occurs if you execute the **display interface ?** command after the **ifmgr** process is rebooted.
- Remarks: None.

201905210837

- Symptom: The switch is stuck at the **ifmgr** process during startup.
- Condition: This symptom occurs if the following operations have been performed:
 - a. Configure a large number of subinterfaces and IBGP neighbors.
 - b. Configure BGP NSR.
 - c. Save the configuration.
 - d. Reboot the switch. The switch selects the .mdb configuration file to load at startup.
- Remarks: None.

201905200469

- Symptom: The switch acting as a subordinate spine node cannot obtain the IP addresses of loopback interfaces on leaf nodes.
- Condition: This symptom occurs if the connected leaf nodes cannot be accessed.
- Remarks: None.

201905200440

- Symptom: The CPU usage of the **comsh** process is high.
- Condition: This symptom occurs if the following operations have been performed:
 - a. Set the locking key to **Ctrl+A** for a VTY line.
 - b. Telnet to the switch, and press **Ctrl+A** to lock the current user line.
 - c. Disconnect the Telnet connection.
 - d. View the CPU usage and process information.
- Remarks: None.

201906280987

- Symptom: Communication exceptions occur.
- Condition: This symptom might occur if the following conditions exist:
 - LLDP is configured.
 - The intrusion protection feature is configured, with the intrusion protection action as **disableport-temporarily**.
 - The number of learned MAC addresses reaches the maximum number of secure MAC addresses.
- Remarks: None.

201906251191

- Symptom: The **vtep access port** command configuration is lost on an aggregate interface.
- Condition: This symptom occurs after you modify the system name of the server after the switch comes online through automatic deployment.

- Remarks: None.

201907231134

- Symptom: The session timeout information still exists in the **display dot1x connection** command output after the server deletes the Session-Timeout attribute during an 802.1X reauthentication.
- Condition: This symptom occurs if the server assigns the Session-Timeout attribute to an 802.1X user during the first authentication and then deletes the Session-Timeout attribute during an 802.1X reauthentication.
- Remarks: None.

201905200477

- Symptom: The switch does not respond to commands issued through NETCONF.
- Condition: This symptom occurs if the commands are issued quickly without delays between them.
- Remarks: None.

201907011015

- Symptom: After a server migrates from a DR interface to a non-DR interface on the primary member device of a DR system, the server cannot ping the gateway attached to the DR system.
- Condition: This symptom might occur if a server migrates from a DR interface to a non-DR interface on the primary member device of a DR system.
- Remarks: None.

201907020289

- Symptom: A user fails MAC authentication on an interface if its MAC address has been learned by another interface of the switch.
- Condition: This symptom might occur if a MAC authentication user accesses an interface and its MAC address has been learned by another interface of the switch.
- Remarks: None.

201907110046

- Symptom: Layer 3 packets of a VSI fail to be forwarded out of a different VSI on another device.
- Condition: This symptom occurs in a distributed VXLAN IP gateway deployment.
- Remarks: None.

201905210842

- Symptom: Multiple Telnet users exist and cannot be deleted after certain operations are performed.
- Condition: This symptom might occur if the following operations are performed:
 - Telnet to the switch from multiple terminals.
 - On each terminal, execute the **telnet 127.0.0.1** command multiple times and press Ctrl + K.
 - Execute the **display users** command on the switch.
- Remarks: None.

201907031120

- Symptom: A VM in an EVPN network cannot receive the NA packets sent by the gateway.
- Condition: This symptom might occur if the following conditions exist:
 - The VM accesses the EVPN network through an Ethernet service instance in VLAN access mode.

- The SVLAN ID that the Ethernet service instance matches is the PVID of the interface that hosts the Ethernet service instance.
 - The VM sends untagged NS requests.
- Remarks: None.

201906030510

- Symptom: Both CLI-configured and controller-issued VSIs exist on the switch. When the running configuration is saved and then the switch is rebooted, some CLI-configured VSIs are lost.
- Condition: This symptom might occur if VSIs are configured both at the CLI and from the controller and the switch is rebooted.
- Remarks: None.

201908090682

- Symptom: Some VMs attached to a DRNI system have traffic loss.
- Condition: This symptom might occur if the RARP protocol packets sent by VMs are forwarded through the IPL.
- Remarks: None.

201908260085

- Symptom: An NMS cannot access the switch through SNMP after accessing the hh3cTunnelVxlanIfEntry node.
- Condition: This symptom might occur if the switch has VXLAN-DCI tunnels and an NMS accesses the hh3cTunnelVxlanIfEntry node.
- Remarks: None.

201909020125

- Symptom: In a VXLAN network, VMs in the same subnet might fail to ping one another after certain operations are performed.
- Condition: This symptom might occur if the following operations are performed:
 - a. Create an Ethernet service instance on an interface, configure the SVLAN ID match criterion for it, and set the PVID of that interface to be the same as the SVLAN ID in the match criterion.
 - b. Delete the Ethernet service instance.
 - c. Repeat steps a and b.
- Remarks: None.

201908010003

- Symptom: The virtual IP addresses of new VRRP groups cannot be pinged after the number of VRRP groups exceeds 512.
- Condition: This symptom might occur if more than 512 VRRP groups are configured.
- Remarks: None.

201905200485

- Symptom: A peer cannot learn the management IP address of the switch through LLDP after the LLDP process restarts on the switch.
- Condition: This symptom might occur if the LLDP process restarts on the switch.
- Remarks: None.

201908060029

- Symptom: The status of the interface on the switch becomes inconsistent with that of the peer interface on the directly connected device after the peer interface goes down and then immediately comes up.
- Condition: This symptom might occur if the peer interface on the directly connected device goes down and then immediately comes up.
- Remarks: None.

201908280757

- Symptom: Layer 3 traffic forwarding is interrupted.
- Condition: This symptom might occur after you disable packet statistics for the Layer 3 aggregate subinterface by using the **undo traffic-statistic enable** command.
- Remarks: None.

201905160399

- Symptom: The CPU usage keeps at 100% for a long time after a recursion loop occurs.
- Condition: This symptom might occur if the following conditions exist:
 - The device has two BGP routes, route **1** and route **2**. Route **1** has a primary next hop **a** and a backup next hop **b** (specified by using FRR); route **2** has a primary next hop **b** and a backup next hop **a** (specified by using FRR).
 - Both **a** and **b** are on the same network segment as the destination networks of route **1** and route **2**.
 - The interfaces pointing to both **a** and **b** go down within a short period of time. As a result, the device selects the backup next hop for both routes. A recursion loop occurs.
- Remarks: None.

202103111262

- Symptom: After the **undo port outbound-vlan-tag enable** command is executed to disable the function of setting VLAN tags for sent protocol packets on an interface, BFD packets still carry VLAN tags. When the physical state of an aggregate interface is down, BFD sessions on the interface are still up.
- Condition: This symptom occurs if the following operations are performed:
 - a. Use the **port outbound-vlan-tag enable** command to enable the function of setting VLAN tags for sent protocol packets on an interface. The BFD sessions are up.
 - b. Disable the function of setting VLAN tags for sent protocol packets on the interface. The BFD sessions are still up.
- Remarks: None.

202103161004

- Symptom: Some tunnel source addresses are not the configured addresses, resulting in traffic forwarding failure.
- Condition: This symptom might occur when a large number of tunnels exist and you modify the tunnel source addresses multiple times.
- Remarks: None.

202104160461

- Symptom: The device might reboot unexpectedly.
- Condition: This symptom occurs if the following operations are performed:
 - Create more than 2000 VLAN interfaces, which are all up.
 - Execute the **info-center enable** command to enable the information center function.

- Unplug and plug subcards to make all VLAN interfaces go down and then come up within a short period of time.
- Remarks: None.

202202140950

- Symptom: On an EVPN+DRNI network, a MAC address entry is learned incorrectly.
- Condition: This symptom occurs if an AC is configured to match untagged packets, and then the MAC address is moved from a DR interface to a single-homed AC interface.
- Remarks: None.

202103151655

- Symptom: Resilient load sharing does not take effect.
- Condition: This symptom might occur if the member ports of an aggregate interface go down after the **link-aggregation load-sharing mode resilient** command is executed on it.
- Remarks: None.

Resolved problems in R2702

201905200485/201901090410

- Symptom: On the IRF fabric, the management address fails to be displayed in the LLDP information received from the neighboring devices.
- Condition: This symptom might occur if the following conditions exist:
 - VLAN interfaces are created on the IRF fabric and IP addresses are assigned to the interfaces.
 - An IRF subordinate device reboots.

201812060001

- Symptom: The XMLCFGD process creates a core file unexpectedly.
- Condition: This symptom might occur if a NETCONF connection is established to the device to manage the device and NETCONF is used to reboot the device.

201809290321

- Symptom: On a DRNI network, a device reboots because of memory exhaustion.
- Condition: This symptom might occur if the following conditions exist:
 - The keepalive timeout timer on the secondary DR member device is set to the maximum value.
 - A configuration rollback is performed on the primary DR member device to cancel the DRNI configuration and then another configuration rollback is performed to recover the DRNI configuration.

201902010798

- Symptom: A device management user fails to obtain another user role by using the **super** command.
- Condition: This symptom might occur if the device management user logs in to the device after passing HWTACACS authentication and executes the **super** command to obtain another user role.

201904010489

- Symptom: The device fails to forward traffic correctly.

- Condition: This symptom might occur if a loop exists on the device, which causes the ARP table to update repeatedly and then causes FIB table update failure.

201903211294

- Symptom: The device reboots unexpectedly.
- Condition: This symptom might occur if the control plane deploys entries that contain unassigned IP addresses to the data plane on a control-/data-plane separated network.

201807190673

- Symptom: The ofcd process fails because of exception.
- Condition: This symptom might occur if the established OpenFlow tunnel is attacked by exception OpenFlow packets in which the length of the protocol header field is 0.

201809110564

- Symptom: The cp process still remains on the device after the connection to the controller is terminated.
- Condition: This symptom might occur if the controller deploys the **save** command through NETCONF to save the running configuration and then terminates the connection to the device.

201811060548

- Symptom: The CPU usage rises rapidly during inter-VPN traffic forwarding.
- Condition: This symptom might occur if BGP redirects direct routes between multiple VPN instances.

201809200079

- Symptom: The RADIUS server fails to assign an authorization VLAN name to a user after the user passes authentication.
- Condition: This symptom might occur if the authorization VLAN name is in the format of \000XXXXX\000.

201904010490

- Symptom: The device reboots unexpectedly.
- Condition: This symptom might occur if ARP entries are deleted when SNMP is walking the ARP table.

201904020841

- Symptom: The device reboots unexpectedly.
- Condition: This symptom might occur if TCP MSS is set on a subinterface and the subinterface is repeatedly deleted and created when SLB traffic is forwarded.

201807300378/201905090714

- Symptom: A memory leak occurs on the SNMP process.
- Condition: This symptom occurs if the following conditions exist:
 - a. SNMP notifications for system logs are disabled.
 - b. The NMS walks the SYSLOG-MSG-MIB to obtain data.

201811070579

- Symptom: The lauthd process creates a core file unexpectedly.
- Condition: This symptom might occur if the **local-user-export class network guest url b** command is executed consecutively several times.

201811060248

- Symptom: The IMC server forcibly logs out a portal user after the user passes portal authentication.
- Condition: This symptom might occur if the portal authentication server runs IMC PLAT 7.3 and security policy confirmation (such as ACL and VLAN) is deployed on the IMC server.

201810230548/201809120806

- Symptom: A memory leakage occurs on a subordinate device in an IRF fabric.
- Condition: This symptom might occur if portal users that obtain IP addresses through DHCP carries Option 82 or Option 18 when they come online.

201809200058

- Symptom: The Aaad process on an IRF fabric creates a core file unexpectedly.
- Condition: This symptom might occur if the following conditions exist:
 - A large number of IPoE users come online through the IRF fabric.
 - Master/subordinate switchover repeatedly takes place.
 - The AAA process reboots repeatedly.

201812070009/201812061078

- Symptom: Specific UDP packets get lost during forwarding.
- Condition: This symptom might occur if a UDP packet has the following characteristics:
 - The packet is a fragment packet.
 - The packet carries MPLS labels.
 - The third and fourth bytes in the IP header of non-first fragment packets is 0D AF.

201811060034

- Symptom: An IPsec SA is established between the device and the peer device through IKEv2 negotiation and the security protocol is ESP. IPsec protocol packets from the peer device are discarded because the packet length exceeds the port MTU.
- Condition: This symptom might occur if TFC padding is enabled and IPsec packet fragmentation is disabled on the peer device.

201903211236

- Symptom: The CLI of a device in an IRF fabric gets stuck and no commands can be input.
- Condition: This symptom might occur if a large number of tunnels flap and IRF master/subordinate switchover repeatedly takes place.

201902020055

- Symptom: IS-IS neighbor relationship cannot be established.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure the network type as P2P and enable IS-IS on an interface.
 - b. Reboot the device.

201904020277

- Symptom: ARP entries become blackhole entries, and packets are lost.
- Condition: This symptom occurs if the following operations are performed:
 - a. Multiple Layer 2 aggregation groups exist in the network, and loops exist in some aggregation groups.
 - b. Enable ARP active acknowledgement.

- c. Configure static routes on a Layer 3 interface. Shut down and then bring up the Layer 3 interface, or MAC address moves occur on the Layer 3 interface.

201902020232

- Symptom: The master IRF member device might reboot unexpectedly at a low probability.
- Condition: This symptom occurs if the following operations are performed:
 - a. Set a small idle timeout value for TCP connections.
 - b. Initiate a large number of TCP connections for services using TCP (for example, BGP and HTTP) on the local end.

201811060022

- Symptom: The memory leaks for the IPFS module.
- Condition: This symptom occurs if the following conditions exist:
 - A large amount of traffic with varying quintuples is forwarded by software.
 - The fast forwarding entries age out.

201902020140

- Symptom: After the TCP client connection is closed, the memory leaks.
- Condition: This symptom occurs if the following operations are performed:
 - a. The client sends a large amount of data to the server. The server cannot process so much data, so the server responds with Zero Window.
 - b. The client starts the persist timer after receiving Zero Window.
 - c. The client actively closes the connection.

201902020187

- Symptom: The CPU usage might be high at a low probability.
- Condition: This symptom occurs if a large number of packets are transmitted when a user logs in through nested Telnet.

201812070478

- Symptom: An interface on a subordinate IRF member device cannot join a voice VLAN again after leaving the voice VLAN.
- Condition: This symptom occurs if the following operations are performed:
 - a. Enable LLDP on an interface on a subordinate IRF member device, and configure a voice VLAN on the interface. Connect the interface to a voice device supporting LLDP/CDP.
 - b. Establish or disconnect the LLDP neighbor relationship on the subordinate IRF member device.

201811060177

- Symptom: After an IP phone successfully comes online, the gateway cannot ping the IP phone for a period of time.
- Condition: This symptom occurs if the following operations are performed:
 - a. Connect an interface to a Cisco IP phone, enable CDP-compatible LLDP on the interface, and assign the IP phone to a voice VLAN.
 - b. The interface repeatedly comes up and goes down.

201811060399

- Symptom: A DHCP client cannot obtain an IP address.
- Condition: This symptom occurs if the device acts as a DHCP sever, multiple address pools are configured, and some address pools are configured with address ranges for dynamic allocation by using the **address range** command.

201812060884

- Symptom: The XMLCFGD process exits exceptionally.
- Condition: This symptom occurs if the following operations are performed:
 - a. The device acts as a DHCP Sever. In a DHCP address pool, configure more than 13 static IP address bindings.
 - b. Use SoapUI to get the data of the DHCP/DHCPStatic table.

201810290644

- Symptom: During auto upgrade, the **using tengige** command is mistakenly executed. As a result, the comsh process becomes abnormal, and related interfaces disappear.
- Condition: This symptom occurs because the **using tengige** command is mistakenly executed during the configuration recovery process. On the device, the **using tengige** command takes effect in real time, but the configuration file incorrectly contains the command.

201903290556

- Symptom: Interface flapping causes the CPU usage to reach 100%.
- Condition: This symptom occurs if the following operations are performed:
 - a. Multiple routes of BGP neighbors are configured with FRR. The active and backup next hops of FRR are reverse for two routes (for example, the active and backup next hops of route A are 1 and 2, and the active and backup next hops of route B are 2 and 1), and the next hops 1 and 2 are in the network segments of routes A and B.
 - b. Shut down the interfaces corresponding to the two next hops in sequence.

201903290558

- Symptom: When the spanning tree mode is switched to PVST, the device will be stuck for a period of time.
- Condition: This symptom occurs if a large number of VLANs and interfaces exist on the device and the spanning tree mode is switched to PVST.

201811060535

- Symptom: When an interface card is unplugged and plugged, the aggregate interface creation event on the interface card is not reported. As a result, the aggregate interface on the interface card is not set to the drive, and the aggregate interface member ports cannot forward traffic.
- Condition: This symptom occurs because the interface management module does not report the aggregate interface creation event during the startup process when an interface card is plugged.
- Occurrence probability: This symptom occurs only when interface events are not reported. In an environment, there are a large number of interface events. In a complicated environment, the occurrence probability is high. In a test environment, the occurrence probability is low.

201807060250

- Symptom: Some traffic is broadcast on a DR interface.
- Condition: This symptom occurs if an aggregate interface leaves and then joins a DR group and continuously receives traffic.

201903110087

- Symptom: The BFD session on a Layer 3 aggregate interface flaps.
- Condition: This symptom occurs if the following operations have been performed:
 - a. Configure a Layer 3 aggregate interface with member ports on different cards, enable BFD for OSPF, and use MD5 authentication for BFD control packets.
 - b. Remove a member port from the Layer 3 aggregation group and then add it back to the aggregation group.

201806040598

- Symptom: The secure MAC address entry is not removed from the **display mac-address** command after a user goes offline.
- Condition: This symptom occurs if port security is configured and the user goes offline after passing authentication.

201701100257

- Symptom: Traffic detection fails in a Fabric Director scenario.
- Condition: This symptom occurs if a QoS policy is issued multiple times.

201806070741

- Symptom: The **remark dscp** command issued by OpenFlow does not take effect.
- Condition: This symptom occurs if the Output action is issued by OpenFlow at the same time.

201904020301

- Symptom: The relevant MAC address entry is not removed from the **display mac-address** command after an 802.1X user moves to a different VLAN on the same port.
- Condition: This symptom occurs if an 802.1X user moves to a different VLAN on the same port.

201904020262

- Symptom: In an EVPN distributed relay environment, the interface where a single-armed AC is configured cannot forward packets.
- Condition: This symptom occurs if the IPP interface setting is cancelled and then restored for a tunnel interface .

201904110239

- Symptom: A DR system fails to be established.
- Condition: This symptom occurs if a manually created tunnel interface is used as the IPL.

201903150058

- Symptom: In a DRNI network, the DR interface of the secondary DR device is still up after the IPP interface is brought down.
- Condition: This symptom occurs if the secondary DR device is in DRNI MAD DOWN state.

201903210720

- Symptom: In an EVPN distributed relay environment, the DR system sends out multiple copies of unknown unicast packets.
- Condition: This symptom might occur if the following operations have been performed:
 - a. Use a VXLAN tunnel as the IPL and reboot the DR system.
 - b. Receive unknown unicast packets from the remote AC.

201812060999

- Symptom: In a DRNI network, the DR interface is set to DRNI DOWN state.
- Condition: This symptom might occur if the IPP interface flaps.

201903080004/201903070270

- Symptom: In an MPLS network, a P device drops packets continuously.
- Condition: This symptom might occur if the link between the P device and another P device or a PE device flaps for a long time more than once.

201805040745

- Symptom: In a multiple VSC environment, the device cannot connect to the primary VSC.
- Condition: This symptom might occur if the OVSDB process is restarted.

201902140542

- Symptom: In an EVPN distributed relay environment, the IPL cannot work correctly.
- Condition: This symptom might occur if you configure VLAN-based VXLAN assignment and then configure EVPN distributed relay.

201810300310

- Symptom: The management Ethernet port goes down in an IRF fabric.
- Condition: This symptom might occur after a master/subordinate switchover is performed.

201711070993

- Symptom: In a VXLAN network, VMs in different network segments cannot communicate.
- Condition: This symptom occurs if a VXLAN gateway group is used as the gateway.

201805020138/201805020139

- Symptom: An additional coldStart log is printed every time the switch sends a trap.
- Condition: This symptom occurs after the switch reboots.

201904020313

- Symptom: A user can join and leave the multicast group without passing authentication.
- Condition: This symptom occurs if both MLD and IPv6 portal authentication are configured on the VLAN interface.

201903180860

- Symptom: A serial port hangs in a DRNI network.
- Condition: This symptom might occur if the following operations have been performed:
 - a. Enable and disable configuration consistency check repeatedly.
 - b. Execute the **display drni consistency type2 global** command.

201810100474

- Symptom: ICMPv6 packets are counted into the **IP-other** protocol type.
- Condition: This symptom occurs when the switch receives ICMPv6 packets.

201811090192

- Symptom: The MAC address entry is not removed from the **display mac-address** command after a MAC authentication user goes offline.
- Condition: This symptom occurs if the MAC authentication user comes online and then goes offline.

201812110026

- Symptom: In an EVPN network, an access port sends packets with VLAN tags.
- Condition: This symptom might occur if two route reflectors are used and link switchover between them has occurred.

201904030323

- Symptom: The remote host has the TCP timestamps vulnerability.
- Condition: This symptom occurs if the host implements RFC 1323.

201812061014

- Symptom: HPE Comware 7 stored and reflected XSS Vulnerability
- Condition: An xss reflected in the web portal of the appliance HP Comware switch 7.1.045. Attackers can exploit this issue to open a web browser and log in to the application using valid or not credentials.

201812050851

- Symptom: Files in the flash might fail to be deleted at a low probability.
- Condition: This symptom occurs if multiple consoles operate the device simultaneously.

201903140269/201904020861

- Symptom: After the operating mode of a device is switched from L3GW to L2GW, the L3VNI configuration remains.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure the device to operate in L3GW mode, and configure L3VNIs.
 - b. Configure the device to operate in L2GW mode, save the configuration, and reboot the device.

201903280399

- Symptom: When EVPN and DRNI are used together on the switch, frequent tunnel interface flapping might cause traffic interruption.
- Condition: This symptom might occur if frequent tunnel interface flapping occurs.

201812280633

- Symptom: The startup configuration file on an IRF subordinate device is lost after a user logs out of the IRF fabric while the IRF fabric is saving the running configuration to the startup configuration file for the user.
- Conditions: This symptom occurs if the user logs out of the IRF fabric while the IRF fabric is saving the running configuration to the startup configuration file for the user.

Resolved problems in F2701

None.

Resolved problems in R2612P02

201809120302

- Symptom: Multiple copies of packets mirrored by Layer 2 remote port mirroring are received.
- Condition: This symptom occurs if the following operations are performed:
 - a. Create multiple mirroring groups, and assign ports to mirroring groups.
 - b. Configure reflector ports for remote mirroring groups.

201809050657/201805120132

- Symptom: In an EVPN network, the VSI IPv6 gateway cannot be pinged.
- Condition: This symptom occurs if you ping the VSI IPv6 gateway bound to a VXLAN through the AC link of the VXLAN.

201809050319/201808230872

- Symptom: After NETCONF is used to deploy the BFD-related configuration, the BFD process fails to start.
- Condition: This symptom occurs if NETCONF is used to deploy the BFD configuration.

201809050305

- Symptom: When an IPL fails, the corresponding Layer 3 interfaces cannot properly learn ARP entries. As a result, traffic is interrupted.
- Condition: This symptom occurs if the following operations are performed:
 - a. In a DRNI network, configure the same MAC address for the VLAN interfaces of the VLANs to which the DR interfaces of the IPL belong.
 - b. Shut down the IPL.

201809040359/201809030027/201809030023

- Symptom: After an IRF master/subordinate switchover, the AC configuration on the device might be deleted and the VM traffic cannot be forwarded at a low probability.
- Condition: This symptom occurs if the following operations are performed:
 - a. On an IRF fabric, the controller automatically deploys the VXLAN function.
 - b. Reboot the master IRF member device.

201809250358/201809250374/201809250364

- Symptom: Traffic cannot be forwarded in a VPLS network.
- Condition: This symptom occurs if the following operations are performed:
 - a. Create an aggregate interface, and configure the aggregate interface as an AC.
 - b. Remove member ports from the corresponding aggregation group and then assign these member ports to the aggregation group.

Resolved problems in R2612P01

201807270157/201806210622

- Symptom: When you use Director to replace the master spine device, the leaf device configuration changes.
- Condition: This symptom occurs if the automated VCF fabric deployment function is used to enable the device to cooperate with Director and implement automated configurations.

201807270712/201807270721/201807270711

- Symptom: After a master/subordinate switchover, an IRF fabric sends redundant RSCN packets to servers.
- Condition: This symptom occurs if the following operations are performed:
 - a. In an FCoE network, enable hardware zoning and configure RSCN on an IRF fabric.
 - b. Reboot the master IRF member device.

201807170431/201807270554

- Symptom: The PCB version for an LSWM124XG2QL interface module is displayed incorrectly.
- Condition: This symptom occurs if the device has an LSWM124XG2QL interface module installed and the **display device verbose** command is executed to view the PCB version information.

201807270202/201807070110/201807070109

- Symptom: The memory leaks after the device configured with a large number of MPLS TE tunnels runs for a period of time.
- Condition: This symptom occurs if the device is configured with a large number of MPLS TE tunnels and the RSVP summary refresh feature and the reliable RSVP message delivery feature are enabled.

201807270185/201807190739/201807190721

- Symptom: CRC error packets exist on an FC interface.
- Condition: This symptom occurs if an FC interface on an LSWM124XG2QFC interface module has an HP 16Gb FC/10GbE 100m SFP+ XCVR transceiver module installed.

201808060501/201808060502/201808060503

- Symptom: The controller might fail to deploy flow entries to the subordinate IRF member devices.
- Condition: This symptom occurs if the following operations are performed:
 - a. An IRF fabric acts as an OpenFlow switch and establishes a secure channel with the controller.
 - b. The controller deploys flow entries to the subordinate IRF member devices.

201808160237/201805220100

- Symptom: An IRF fabric might fail to forward Layer 2 packets on a multichassis Layer 2 aggregate interface. VRRP flapping occurs.
- Condition: This symptom might occur when the IRF fabric is configured with both VRRP and SPBM.

201807270142/201806200386/201805300594

- Symptom: After the DR interface comes up, it will go down and then come up once.
- Condition: This symptom occurs if you view the DR interface status after the IPL comes up.

201807270145/201806250510/201803190222

- Symptom: A client cannot join a multicast group.
- Condition: This symptom occurs if the client comes online through portal and requests to join the multicast group in a multicast network.

201807270155/201806250584/201806060558

- Symptom: An interface does not respond after receiving ARP requests.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure an Ethernet service instance on an aggregate interface, and configure the **encapsulation default** command on the Ethernet service instance.
 - b. Configure the **undo encapsulation** command on the Ethernet service instance to restore the default.
 - c. The aggregate interface receives ARP requests.

201807270161/201806120577/201806270423/201806120577

- Symptom: After the reload delay timer set for a DR device expires, the DR device role is still None.
- Condition: This symptom occurs if the following operations are performed:
 - a. Execute the **drni auto-recovery reload-delay delay-value** command to enable DR system auto-recovery and set the reload delay timer.
 - b. Configure both the IPP and keepalive link to be down.

- c. Save the configuration and reboot the DR device.

201807270168/201806270402/201806070375/201806070389

- Symptom: When the **display drni role** command is used to display DR role information on the secondary DR device, the **Effective role** field displays **Primary**.
- Condition: This symptom occurs if the IPP is repeatedly shut down and brought up in a DRNI network.

201807270130/201807030034/201806290366/201806290360

- Symptom: After the whole IRF fabric is rebooted, SNMP obtains an incorrect value for the snmpEngineBoots node.
- Condition: This symptom might occur if the master member device of the IRF fabric changes after the IRF fabric is rebooted.

201807270124/201806270357/201806040701

- Symptom: The chip time is different on the master IRF member device and subordinate IRF member device.
- Condition: This symptom occurs if an IRF fabric is configured with PTP and the chip time on the master IRF member device and subordinate IRF member device is viewed.

201806260327/201807270176/201807060219

- Symptom: A DR system fails after certain operations are performed.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure a tunnel interface as the IPP.
 - b. Configure dynamic tunnels on the DR member devices, and the dynamic tunnels share the destination IP address with the tunnel that acts as the IPL.
 - c. Delete the IPP tunnel interface and reconfigure it.

201807270182/201807030632/201807310533

- Symptom: On the secondary DR member device, a DR interface in DRNI DOWN state is removed from its DR group. After the DR interface is reassigned to the DR group, its state becomes DOWN.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure a Layer 2 aggregate interface as a DR interface and assign it to a DR group on the secondary DR member device.
 - b. Remove the DR interface from its DR group and then reassign it to the DR group when the IPL is down.

201807270188/201806010178/201807030865

- Symptom: On a DR member device, member ports of a DR interface cannot become Selected after the device is rebooted.
- Condition: This symptom might occur if the following operations are performed:
 - a. Execute the **lacp edge-port** command on the DR interface.
 - b. Save the configuration and reboot the DR member device.

201807270193/201807070082/201807070098

- Symptom: RSVP has memory leaks if RSVP authentication fails.
- Condition: This symptom might occur if RSVP authentication fails.

201807270196/201807100205/201807100209

- Symptom: Memory leaks occur if the switch repeatedly generates and deletes a large number of multicast entries.

- Condition: This symptom might occur if the switch repeatedly generates and deletes a large number of multicast entries.

201807060212/201807270199/201807060355

- Symptom: DR member devices might fail to forward Layer 3 traffic after certain operations are performed.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure a VXLAN tunnel interface as the IPP of the DR system.
 - b. Configure the DR member devices to establish dynamic tunnels to external networks.
 - c. Delete the VXLAN tunnel interface.
 - d. Shut down and then bring up the interfaces connected to the external networks.
 - e. Create a VXLAN tunnel interface and configure it as the IPP.

201807270206/201806290774/201807100295

- Symptom: Third-party services, service chain, and PBR are configured on an 5940 switch that acts as a leaf node in a VCF fabric. After the **reset arp all** command is executed, PBR configuration does not take effect.
- Condition: This symptom might occur if the **reset arp all** command is executed on the 5940 switch.

201807270207/201806280646/201806270600

- Symptom: Memory leaks for the OVSDb module. About 50 bytes leak every 10 seconds. If the controller re-deploys the configuration, about 80 bytes leak.
- Condition: This symptom occurs if the device has the OVSDb service enabled, and the data in the OVSDb database are modified after the controller deploys a global table containing the master controller IP to the OVSDb database.

201807170279/201807270285/201804100876

- Symptom: The device name configured for a device by using the **sysname** command does not take effect.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure automated underlay network deployment on the device.
 - b. Use the **sysname** command to modify the device name, save the configuration, and reboot the device.

201807270468/201805220131

- Symptom: After the device runs for a period of time, the MACsec data packets cannot be forwarded.
- Condition: This symptom occurs if the device acts as a MACsec client and establishes a device-oriented MACsec network with a Huawei or Cisco device.

201807270215/201806260106/201806250611

- Symptom: The device might reboot unexpectedly.
- Condition: This symptom occurs if an aggregation group has more than 32 member ports and any member port leaves the aggregation group.

201808210067

- Symptom: If the **undo irf mac-address persistent** command is executed on an IRF fabric configured with VXLAN, overlay traffic forwarding fails after an IRF master/subordinate switchover.

- Condition: This symptom might occur if the **undo irf mac-address persistent** command is executed on an IRF fabric configured with VXLAN, and an IRF master/subordinate switchover occurs.

Resolved problems in R2612

201804100157/201804100168/201804100163

- Symptom: Traffic cannot be forwarded between virtual machines.
- Condition: This symptom occurs if the following conditions exist:
 - a. VTEPs and remote VTEPs are connected through HP flood proxy servers.
 - b. Enable flood proxy on multiple tunnels, and bind the tunnels to VSIs.
 - c. Enable ARP flood suppression on the device.
 - d. The local proxy tunnel is different from the tunnel selected by HP servers.

201805120143/201805120129

- Symptom: Auto-RP listening does not take effect.
- Condition: This symptom might occur if the following operations are performed:
 - a. Enable Auto-RP listening on the device.
 - b. Configure a Layer 2 aggregate interface as a trunk port and assign it to a VLAN.
 - c. Enable PIM-SM on the VLAN interface.

201805110585

- Symptom: On a DRNI+STP network, traffic interruption occurs after the IPL goes down and then comes up.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure an aggregate interface on each DR member device as the IPP.
 - b. Reboot a DR member device so that the DR member devices are assigned new roles.

201805110581/201803270029

- Symptom: On a DRNI+STP network where the DR system operates correctly, it takes DR interfaces ten minutes to come up after they are set to DRNI MAD DOWN state.
- Condition: This symptom might occur if a DR member device reboots and then the IPP goes down.

201805110456

- Symptom: The ovsdb-server process exits unexpectedly.
- Condition: This symptom might occur after a VTEP is enabled with the OVSDB server feature and establishes an OVSDB connection with the controller.

201805090685

- Symptom: An IRF subordinate device reboots unexpectedly after the **display interface** command is executed on the IRF fabric.
- Condition: This symptom might occur if Layer 3 Ethernet subinterfaces are created on the IRF fabric.

201805070301

- Symptom: The OVSDB connection to the controller is disconnected after a length of time since a VTEP has been enabled with the OVSDB server feature and established an OVSDB connection to the controller.

- Condition: This symptom might occur after a length of time since a VTEP has been enabled with the OVSDDB server feature and established an OVSDDB connection to the controller.

201805050184

- Symptom: The device fails to set the VXLAN hardware resource mode.
- Condition: This symptom might occur if the following operations:
 - a. Set the VXLAN hardware resource mode.
 - b. Save the running configuration and reboot the device.
 - c. Use the **display hardware-resource** command to display the VXLAN hardware resource mode. The displayed hardware resource mode is not the specified one.

201805100244

- Symptom: The remote fault signal detection feature, which is supported only on fiber ports, can be enabled on copper ports.
- Condition: This symptom might occur if the **link-fault-signal enable** command is executed on copper ports.

201805090323

- Symptom: The system prompts unsupported operation if the speed of a 100-GE interface is repeatedly changed between 100000 Mbps and 10000 Mbps.
- Condition: This symptom might occur if the speed of a 100-GE interface is repeatedly changed between 100000 Mbps and 10000 Mbps by using the **speed 100000** and **speed 10000** commands.

201805040458

- Symptom: The memory of the QACL module slowly leaks.
- Condition: This symptom occurs if actions in traffic behaviors are dynamically modified repeatedly.

201805020139

- Symptom: The device prints coldStart traps unexpectedly when printing port security traps.
- Condition: This symptom occurs when the device is rebooted and prints port security traps.

201805020133

- Symptom: When the device learns secure MAC address entries, it prints the same traps for twice.
- Condition: This symptom occurs if the device has port security enabled and is configured with secure MAC address entries.

201804270553

- Symptom: The VRRP advertisement packets received by the switch match a PBR policy. As a result, the VRRP advertisement packets fail to be transparently transmitted.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure a PBR policy on a VLAN interface of the device.
 - b. The VLAN interface receives and sends VRRP advertisement packets.

201804250026

- Symptom: The Connect Retry timer times out. As a result, BGP might flap.
- Condition: This symptom occurs if the following operations are performed:
 - a. On an IRF fabric, configure BGP NSR.
 - b. Reboot the device after the device has run for a long period of time.

201804170805

- Symptom: An interface fails to join an aggregation group.
- Condition: This symptom occurs if the following operations are performed:
 - a. Execute the **vtep access port** command to specify a site-facing interface as a VTEP access port.
 - b. Create an aggregation group, and assign the interface to the aggregation group.

201804170540

- Symptom: Failed to read the necBgpEvpnPAtrRouteType node.
- Condition: This symptom occurs if you read the necBgpEvpnPAtrRouteType node through MIB.

201804160611

- Symptom: When the TTL in IPv6 BGP protocol packets is 1, the packets mistakenly match an ACL used for matching IPv6 packets with TTL as 1. As a result, the link flaps.
- Condition: This symptom occurs if IPv6 BGP protocol packets with TTL as 1 are received.

201804120615

- Symptom: A user cannot log in to the device by using NETCONF after certain operations when password control is enabled.
- Condition: This symptom occurs if the following operations are performed:
 - a. Enable password control on the device.
 - b. Repeatedly establish and delete sessions, and perform active/standby process switchover.
 - c. Log in to the device by using NETCONF.

201804120137

- Symptom: In a DRNI network, MAC address entries fail to be synchronized between the primary and secondary devices.
- Condition: This symptom occurs if the following operations are performed:
 - a. In a DRNI, execute the **shutdown** and **undo shutdown** commands on the IPP.
 - b. The device receives a large number of Layer 2 packets with changing source MAC addresses.

201803150880

- Symptom: VLAN-based VXLAN assignment configuration cannot be restored by using an .mdb binary file.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure a large number of VSIs and enable VLAN-based VXLAN assignment.
 - b. Save the configuration, reboot the switch, and use an .mdb binary file to restore the configuration.

201802280277

- Symptom: The controller cannot discover the site-facing interfaces configured by using **vtep access port** if the switch uses Chinese GB2312 characters as the sysname.
- Condition: This symptom might occur if the sysname of the switch contains Chinese GB2312 characters.

201805150032/201712060462/201712060449

- Symptom: The switch reboots unexpectedly.
- Condition: This symptom occurs if the debugging command is used to disable the linkscan for interfaces.

201805100905/201805100908

- Symptom: The switch acts as a VXLAN VTEP, and an Ethernet service instance that uses the **encapsulation default** criterion is configured on an aggregate interface. After the aggregate interface is shut down and then brought up, traffic received on the Ethernet service instance cannot be forwarded correctly.
- Condition: This symptom might occur if an Ethernet service instance that uses the **encapsulation default** criterion is configured on an aggregate interface, and the aggregate interface is shut down and then brought up.

201804240046/201802240168/201709010504

- Symptom: ACLs might remain at a low probability after certain operations.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure a routing policy, and specify the next hop of the routing policy as a GRE tunnel interface.
 - b. Modify the source IP address of the GRE tunnel.

201804200600/201804260132

- Symptom: MPLS is configured on the switch, and an aggregate interface can reach the remote peer through multiple ECMP routes. If the aggregate interface is configured to use the source IP address and source port for load sharing, traffic on the aggregate interface is not evenly distributed among the aggregation member ports.
- Condition: This symptom might occur if an aggregate interface that can reach the remote peer through multiple ECMP routes is configured to use the source IP address and source port for load sharing.

201805290161/201805280462

- Symptom: Disabling MAC address learning does not take effect on a Layer 2 aggregate interface.
- Condition: This symptom occurs if the following operations are performed:
 - a. Disable MAC address learning globally.
 - b. In the view of a Layer 2 aggregate interface, execute the **undo mac-address mac-learning enable** command to disable MAC address learning.

201805290049/201805280775

- Symptom: The CLI does not respond after password control is disabled.
- Condition: This symptom occurs if the following operations are performed:
 - a. Enable Password Control on the device. A large number of invalid NETCONF users log in to the device.
 - b. Disable password control.

201805250503/201805250377

- Symptom: Some ACL resources remain.
- Condition: This symptom occurs if the following operations are performed:
 - a. The switch operates in FCF mode and connects to multiple nodes.
 - b. Modify the bridge MAC address of the switch.

201805240699/201805220499

- Symptom: The device prints deadlock logs when the **step** command is used to modify the rule numbering step for an ACL.
- Condition: This symptom occurs if the following operations are performed:

- a. Configure a PBR policy on the device, and configure rules for the ACL that the PBR policy uses.
- b. Apply the PBR policy to packets that an interface forwards.
- c. Enter the view of the ACL, and use the **step** command to set the rule numbering step.

201805240599/201805150488

- Symptom: OpenFlow is disconnected from the controller.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure OpenFlow on the device and establish a connection to the controller.
 - b. The interface corresponding to the AC is frequently shut down and brought up.

201805310080/201805310084/201805310093

- Symptom: The broadcast packets received on a member port of an aggregation group might be broadcast out of other member ports of the aggregation group.
- Condition: This symptom occurs if the following operations are performed:
 - a. Assign local ports to an aggregation group. Delete the aggregation group. Restore the default settings for member ports, and then assign these ports to the aggregation group.
 - b. Execute the **shutdown** and **undo shutdown** command sequence on the aggregation group member ports.
 - c. Switch the mode of the aggregation group to dynamic or static.
 - d. The local device is an STP root bridge. An interface on the peer device repeatedly flaps, and the peer device sends TCN BPDUs.

201803150151

- Symptom: A VLAN interface is associated with a VPN instance, and an Ethernet subinterface that uses a subinterface number the same as the VLAN interface number is created. After the **traffic-statistic enable** command is configured on the Ethernet subinterface, the VLAN interface cannot forward traffic.
- Condition: This symptom occurs if the following operations are performed:
 - a. Associate a VLAN interface with a VPN instance.
 - b. Create an Ethernet subinterface. The subinterface number is the same as the VLAN interface number.
 - c. Execute **traffic-statistic enable** on the Ethernet subinterface.

201804110841

- Symptom: On an IRF fabric, the connection to the OpenFlow controller is interrupted after certain operations are performed.
- Condition: This symptom occurs if the following operations are performed:
 - a. Set up an IRF fabric.
 - b. Issue VXLAN flow entries through OpenFlow.
 - c. Perform an IRF master/subordinate switchover.
 - d. Deactivate the OpenFlow instance and reactivate it.

201803130127

- Symptom: After the switch reboots, the value of the snmpEngineBoots node becomes incorrect.
- Condition: This symptom occurs if the switch reboots.

201804020574

- Symptom: When packet loss prevention is enabled for OpenFlow forwarding, the **apply qos remarking policy** command cannot be configured on interfaces.

- Condition: This symptom occurs if the **apply qos remarking policy** command is configured after packet loss prevention is enabled for OpenFlow forwarding.

201803150342

- Symptom: When a Python script is used to issue configuration, the switch displays that the flash memory is not readable or writable.
- Condition: This symptom occurs if a Python script is used to issue configuration.

201802240291

- Symptom: Patch installation fails after certain operations are performed.
- Condition: This symptom occurs if the following operations are performed:
 - a. Install the patch by specifying the lower-case file name.
 - b. Uninstall the path by specifying the upper-case file name.
 - c. Reinstall the patch by specifying the lower-case file name.

201711140786/201804100468

- Symptom: The DSCP value of BGP protocol packets becomes 0 after the configuration is restored by using a .cfg configuration file or peer addresses are added to a peer group.
- Condition: This symptom occurs if the configuration is restored by using a .cfg configuration file or peer addresses are added to a peer group.

201803020552

- Symptom: When loop detection is enabled and the loop protection action is set to block, the switch cannot block outgoing packets on a looped interface.
- Condition: This symptom occurs if loop detection is enabled and the loop protection action is set to block.

201803020709

- Symptom: When the OpenFlow controller issues changes of the **vsi interface** or **tunnel interface** setting to the switch, error code OFPPMFC_BAD_PORT instead of OFPPMFC_BAD_CONFIG is returned.
- Condition: This symptom occurs if the OpenFlow controller issues changes of the **vsi interface** or **tunnel interface** setting to the switch.

201803270454

- Symptom: On an LSWM18CQMSEC(JH957A) card, an interface is split into four breakout interfaces. When PFC is configured on the breakout interfaces, it takes effect only on the breakout interface numbered 1.
- Condition: This symptom occurs if an interface on an LSWM18CQMSEC(JH957A) card is split into four breakout interfaces, and PFC is configured on the breakout interfaces.

201803270319

- Symptom: Two switches are connected by ports on an LSWM18CQMSEC card and an 5930 24-port 10GBASE-T and 2-port QSFP+ with MACsec Module JH182A card. When MACsec is enabled on the switches, spanning tree packets and PFC packets are interrupted.
- Condition: This symptom occurs if two switches are connected by ports on an LSWM18CQMSEC card and an 5930 24-port 10GBASE-T and 2-port QSFP+ with MACsec Module JH182A card, and MACsec is enabled on the switches.

201804030057

- Symptom: MKA cannot be enabled after it is repeatedly enabled and disabled.
- Condition: This symptom occurs if MKA is repeatedly enabled and disabled.

201802050250

- Symptom: The switch reboots unexpectedly when an FC interface connected to a server or storage device is repeatedly brought up and shut down.
- Condition: This symptom occurs if an FC interface connected to a server or storage device is repeatedly brought up and shut down.

201803290757

- Symptom: The management Ethernet interface might not come up if the switch is rebooted after a transceiver module is installed in the interface.
- Condition: This symptom occurs if a transceiver module is installed in the management Ethernet interface, and then the switch is rebooted.

201804120752

- Symptom: After certain operations are performed, an IRF fabric splits, and IRF member devices reboot.
- Condition: This symptom occurs if the following operations are performed:
 - a. Set up an IRF fabric.
 - b. Configure 64 voice VLANs and 128 voice VLAN OUI addresses.
 - c. Execute the **undo voice-vlan security enable** and **voice-vlan security enable** commands in sequence when Layer 2 traffic exists.

201804090777/201804100669

- Symptom: IRF member devices might reboot unexpectedly if they send a large number of protocol packets that contain invalid fields to one another.
- Condition: This symptom occurs if IRF member devices send a large number of protocol packets that contain invalid fields to one another.

201803050843/201803010118

- Symptom: If an 802.1X user migrates multiple times between ports, the switch does not have MAC address information for the user and cannot forward its traffic.
- Condition: This symptom occurs if an 802.1X user migrates multiple times between ports.

201803300346

- Symptom: Incoming packet statistics of a VXLAN tunnel are incorrect.
- Condition: This symptom occurs if one of the following situations exists:
 - The VXLAN of the VXLAN tunnel uses head-end replication, and the VXLAN tunnel receives multicast or broadcast packets encapsulated in VXLAN packets.
 - The VXLAN of the VXLAN tunnel uses tandem replication, and the VXLAN tunnel receives unicast packets encapsulated in VXLAN packets.

201712150323/201803300529

- Symptom: The DR member devices in an EVPN distributed relay system have inconsistent ARP information.
- Condition: This symptom occurs if the DR member devices receive identical ARP packets on both DR interfaces and non-DR interfaces.

201804120474

- Symptom: RSVP is enabled on a large number of subinterfaces on the switch. After the **placement reoptimize** command is executed, the rsvp process is abnormal.
- Condition: This symptom occurs if RSVP is enabled on a large number of subinterfaces, and the **placement reoptimize** command is executed.

201804100591

- Symptom: RSVP is enabled on a large number of subinterfaces on the switch. When configuration rollback is performed, the console does not respond.
- Condition: This symptom occurs if RSVP is enabled on a large number of subinterfaces, and configuration rollback is performed.

201802110557

- Symptom: An RRPP network contains an IRF fabric. After the IRF master is rebooted, it takes 6 to 10 seconds for RRPP to converge.
- Condition: This symptom occurs if an RRPP network contains an IRF fabric, and the IRF master is rebooted for a master/subordinate switchover.

201804030613

- Symptom: The **display ip routing-table all-vpn-instance statistics** command does not display information for all VPN instances.
- Condition: This symptom occurs if the **display ip routing-table all-vpn-instance statistics** command is executed.

201804091123

- Symptom: Multicast VPN traffic fails to be forwarded.
- Condition: This symptom occurs if the multicast tunnel goes down and comes up repeatedly.

201803190291

- Symptom: In an EVPN network, a Layer 3 VSI interface cannot forward the traffic received from the other end of the VXLAN tunnel.
- Condition: This symptom occurs after a MAC address is configured for the VSI interface and then deleted.

201803260605

- Symptom: In a VXLAN network, a member device in the IRF fabric reboots unexpectedly.
- Condition: This symptom occurs when a 40-GE interface on the member device is split into four 10-GE breakout interfaces by using the **using tengige** command.

201803260912

- Symptom: MPLS settings do not take effect on the 10-GE breakout interfaces of a 40-GE interface.
- Condition: None.

201803160523

- Symptom: The switch does not support BiDi transceiver modules TX1310/RX1490.
- Condition: This symptom occurs when the transceiver modules are inserted in 10-GE interfaces.

201711110094

- Symptom: Multiple configuration files with the same name exist, and some configurations are lost.
- Condition: This symptom occurs after you configure the switch and save the configuration.

201803270694

- Symptom: The **openflow shutdown** command does not take effect on ports of a subordinate device in an IRF fabric.
- Condition: None.

201802110517

- Symptom: The message **Failed to obtain the current configuration** is displayed when the **display current-configuration diff** command is executed.
- Condition: This symptom occurs if the following operations are performed:
 - a. Execute the **display current-configuration diff** command when a large amount of configuration exists on the switch.
 - b. Press CTRL_C to interrupt the display.
 - c. Execute the **display current-configuration diff** command again.

201803140514

- Symptom: 64-byte memory leaks occur.
- Condition: This symptom occurs if the local forwarding capability of a PEX is enabled and disabled repeatedly.

201804090914

- Symptom: The switch cannot learn routes imported from VPN 1
- Condition: This symptom occurs if the following conditions exist:
 - a. The export route target of VPN 1 matches the import route target of only VPN 2.
 - b. The routes of VPN 1 are imported into VPN 2.

201803130750

- Symptom: The switch reboots unexpectedly.
- Condition: This symptom occurs if the following operations are performed:
 - a. Execute the **using tengige** command to split a 40-GE interface into four 10-GE breakout interfaces.
 - b. Execute the **priority-flow-control dot1p ingress-buffer dynamic** command on a 10-GE breakout interface to set the dynamic back pressure frame triggering threshold.

201803130428

- Symptom: The switch cannot forward packets larger than 9000 bytes.
- Condition: This symptom occurs after the cut-through forwarding feature is enabled and then disabled.

201803071063/201803260738

- Symptom: The buffer usage values in the **display buffer usage** command output are incorrectly displayed as 0.
- Condition: None.

201803260937

- Symptom: The TxRx-mode generic flow control function is no longer in effect on an Ethernet interface.
- Condition: This symptom occurs after you forcibly enable PFC and then disable it on the Ethernet interface.

201804030190

- Symptom: The **priority-flow-control dot1p ingress-buffer dynamic** command does not take effect after the switch reboots.
- Condition: This symptom occurs if the following operations are performed:
 - a. Execute **priority-flow-control dot1p ingress-buffer dynamic** command.
 - b. Save the configuration and reboot the switch.

201711150701

- Symptom: No alarm message is generated when QoS and ACL resources are exhausted.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure a large number of ACLs, which exhaust QoS and ACL resources.
 - b. Dynamically modify ACL rules.

201804130368

- Symptom: The **dot1x handshake reply enable** command is stuck on a hybrid port, and the switch cannot learn MAC addresses.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure MAC-based access control, and configure an 802.1X guest VSI on the hybrid port. A large number of users come online through the hybrid port.
 - b. Execute the **dot1x handshake reply enable** command repeatedly on the hybrid port.

201803290621

- Symptom: The **flow-interval** command does not take effect on a PEX in an IRF 3.1 system with the switch acting as the parent device.
- Condition: This symptom occurs if the parent device reboots with the .cfg configuration file.

201803290600

- Symptom: The views of some interfaces on a PEX in an IRF 3.1 system cannot be entered from the parent device. The switch acts as the parent device.
- Condition: This symptom occurs after the parent device is rebooted.

201801130260

- Symptom: The FCoE process exits unexpectedly.
- Condition: This symptom occurs if you restart the FCoE process by using the **process restart name fcoed** command when there is a large number of registered nodes in a VSAN.

201802260584

- Symptom: Layer 3 interfaces in an IRF fabric cannot be configured with IP addresses.
- Condition: This symptom might occur if a master/subordinate switchover is performed repeatedly.

201804200180

- Symptom: In an IRF fabric, the traffic statistics for a VSI interface in the **display interface vsi-interface** command output are incorrect.
- Condition: This symptom occurs if the VSI interface has incoming or outgoing traffic.

201804080535

- Symptom: In a VXLAN network, multicast traffic forwarding errors occur.
- Condition: This symptom occurs if the following conditions exist:
 - a. An aggregate interface acts as the outgoing interface of the multicast tunnel.
 - b. The member ports of the aggregate interface change.

201804040585

- Symptom: Memory leaks occur.
- Condition: This symptom occurs if the following operations are performed:
 - a. Enable PIM-DM on a Layer 3 aggregate interface.
 - b. Bring up and shut down the Layer 3 aggregate interface repeatedly.

201803210461

- Symptom: Traffic might fail to be forwarded.
- Condition: This symptom occurs if the following operations are performed in a multicast VPN:
 - a. Remove a port from a service loopback group.
 - b. After the MTunnel interface goes down, assign the port to the service loopback group.

201803200787

- Symptom: The 40-GE interfaces on the device cannot be configured.
- Condition: This symptom occurs if the following operations are performed:
 - a. Enable preprovisioning on the device. The device is preconfigured.
 - b. Disable preprovisioning on the device.

201803200580

- Symptom: In a VXLAN network, a VTEP mistakenly deletes the inner VLAN tag when forwarding ARP packets with double VLAN tags.
- Condition: This symptom occurs if the following operations are performed:
 - a. Enable ARP flood suppression for a VSI.
 - b. The AC receives ARP packets with double VLAN tags, and floods the ARP packets to remote VTEPs through VXLAN tunnels.

201801020039

- Symptom: When NETCONF is used to delete the DRNI keepalive packet parameters, the UDP port-related configurations remain.
- Condition: This symptom occurs if the following operations are performed:
 - a. Use NETCONF to configure the DRNI keepalive packet parameters. Execute the **drni keepalive ipv6 destination *ipv6-address* udp-port *udp-number*** command.
 - b. Use the **undo drni keepalive ipv6** command to restore the default configuration.
 - c. Execute the **drni keepalive ip destination *ipv4-address* udp-port *udp-number*** command again.

201712290801

- Symptom: Failed to disable flooding for a VSI.
- Condition: This symptom occurs if the following operations are performed:
 - a. Execute the **flooding disable all** command in VSI view. A large number vPorts come online through this VSI.
 - b. Delete all vPorts of the VSI.
 - c. The vPorts come online again through the VSI.

201712150602

- Symptom: The **drni auto-recovery reload-delay** command configuration does not take effect.
- Condition: This symptom occurs if the following operations are performed:
 - a. In a DRNI network, configure the **drni auto-recovery reload-delay *delay-value*** command.
 - b. Restart the DR system. Before the reload delay timer times out, display the DR role information.

201712150295

- Symptom: Deleting aggregation group configuration from a device is very slow.
- Condition: This symptom occurs if the following operations are performed:
 - a. In a VCF fabric network of the distributed VXLAN type, an 5940 switch acts as a leaf, and the automated overlay deployment is implemented by using Neutron.
 - b. Create an aggregation group on the device. Repeatedly assign ports to and remove ports from the aggregation group.

201712141008

- Symptom: The issued ACL rules are cleared.
- Condition: This symptom occurs if the following operations are performed:
 - a. Create an ACL, and configure ACL rules reaching the maximum number for the ACL.
 - b. Use this ACL for packet filtering on multiple ports.
 - c. Add a new rule 0 to this ACL.

201712141002

- Symptom: A user fails to come online through a Layer 2 aggregate interface with MAC authentication enabled.
- Condition: This symptom occurs if the following operations are performed:
 - a. Enable MAC authentication on the Layer 2 aggregate interface. Execute the **mac-authentication carry user-ip** command on the Layer 2 aggregate interface to include user IP addresses in MAC authentication requests sent to an IMC server.
 - b. The user comes online through the Layer 2 aggregate interface.

201712140553

- Symptom: IPv6 LSP statistics are incorrect.
Condition: This symptom occurs if the following operations are performed:
 - a. Execute the **interface tunnel tunnel-number mode mpls-te** command on the device to create an MPLS TE tunnel interface.
 - b. Configure an IP address for the tunnel interface.
 - c. View the IPv6 LSP statistics.

201712130910

- Symptom: A user failed to log in to the device through SSH.
Condition: This symptom occurs if the following operations are performed:
 - a. Enable the SSH server on the device.
 - b. Execute the **authentication login** command in ISP domain view to specify the authentication method for login users.
 - c. The user logs in to the device by using SSH and sends accounting requests.

201712140046

- Symptom: ND entries are not updated.
- Condition: This symptom occurs if the following operations are performed:
 - a. Create a link aggregation management VLAN, and assign an IPv6 address to the VLAN interface of the VLAN.
 - b. Configure a member port as the link aggregation management port.

- c. Create ND entries for the link aggregation management VLAN.
- d. Cancel the link aggregation management port configuration, and specify a new port as the link aggregation management port.

201712110120/201711250135

- Symptom: When the **shutdown** and **undo shutdown** commands are configured on an aggregation group member port, the member port cannot leave and join the aggregation group.
- Condition: This symptom occurs if the links between spine nodes and leaf nodes are automatically aggregated through automated underlay deployment in a VFC fabric network of the VLAN type.

201712080634

- Symptom: The state of a USB flash drive is displayed as **Absent** incorrectly.
- Condition: This symptom occurs if a USB flash drive is inserted into the device and the **display device usb** command is executed after the USB flash drive is successfully mounted.

201712070914

- Symptom: Unselected member ports of a static aggregation group do not drop packets.
- Condition: This symptom occurs if the following operations are performed:
 - a. Execute the **link-aggregation lacp traffic-redirect-notification enable** command in system view to enable link aggregation traffic redirection.
 - b. Create a static aggregation group, and execute the **link-aggregation selected-port maximum max-number** command to set the maximum number of Selected ports allowed for the aggregation group.
 - c. Assign member ports to the aggregation group to make the number of member ports exceed the maximum number of Selected ports allowed.
 - d. The aggregation group receives known unicast packets.

201712070187

- Symptom: The AC's input and output traffic statistics that the switch reports to the controller are incorrect.
- Condition: This symptom occurs if the following operations are performed:
 - a. The controller deploys the VXLAN configuration to the switch.
 - b. The AC receives and sends packets properly.
 - c. The switch sends the AC traffic statistics to the controller.

201712060054

- Symptom: In an EVPN network, the physical transport-facing interface of a VXLAN tunnel fails to forward overlay packets.
- Condition: This symptom occurs if an interface first acts as an IRF physical interface, then the interface is switched to a common service interface, and then the interface is configured as the VXLAN transport-facing interface.

201711060761

- Symptom: The flow entries deployed by the controller are different from the actual flow entries on the device.
- Condition: This symptom occurs if the connection between the controller and the switch is disconnected and then re-connected when the OpenFlow controller is deploying a large number of MAC flow entries to the device.

201801170393

- Symptom: When you use NETCONF to get the locally learned MAC addresses, the remote MAC addresses are displayed in the local MAC address learning information.
- Condition: This symptom occurs if the VTEP learns the MAC addresses of VMs in the remote site and then NETCONF is used to get the MAC addresses of VMs in the local site.

201712270280

- Symptom: PFC configured on a 100-GE interface does not take effect. The switch still receives packets carrying the specified priority.
- Condition: This symptom occurs if PFC is enabled on 100-GE interfaces of the local switch and peer device.

201801170403/201711130710

- Symptom: In a VXLAN network, when the flood traffic is forwarded in tandem replication mode, ARP packets are sent to the CPU and dropped.
- Condition: This symptom occurs if ARP packets are received in tandem replication mode.

201801200013

- Symptom: In an IRF 3.1 system, core files might be generated during the software upgrade for PEXs.
- Condition: This symptom occurs if the **boot-loader pex** command is executed on the parent device to upgrade software for PEXs and some PEXs go offline during the upgrade process.

201801170429/201710200637

- Symptom: The ARP packets received from ACs cannot be sent to the controller.
- Condition: This symptom occurs if the controller deploys a flow entry with the apply-actions instruction and ACs receive ARP packets.

201801130350

- Symptom: Memory leaks for the BGP module.
- Condition: This symptom occurs if either of the following operations is performed:
 - In BGP IPv4 unicast address family view, repeatedly execute the **peer route-policy** and **undo peer route-policy** commands to apply a routing policy to routes incoming from or outgoing to a peer or peer group and remove the configuration.
 - Modify the applied routing policy.

201801120755

- Symptom: The switch intermittently generates the level-1 memory threshold alarms, and BGP routes flap.
- Condition: This symptom occurs if a large number of IPsec SAs repeatedly flap in certain conditions.

201801090902

- Symptom: When the original master IRF member switch is rebooted and upgraded, the BGP neighbor is disconnected.
- Condition: This symptom occurs if ISSU master/subordinate switchover is performed when the BGP routes are being withdrawn and updated in an IRF fabric.

201712050263

- Symptom: An interface on the switch does not come up.
- Condition: This symptom occurs if a 1000-Mbps transceiver module is inserted into the interface and then replaced with a 10-Gbps transceiver module on the 5940 switch.

201711170583

- Symptom: A VLAN interface on the switch cannot forward Layer 3 traffic.
- Condition: This symptom occurs if the VLAN interface is configured with the **arp mode uni** command.

201711140594

- Symptom: When 802.1X is configured on a Layer 2 aggregate interface on a VTEP, Ethernet service instances fail to be dynamically created. As a result, users cannot join the guest VSI.
- Condition: This symptom occurs if the following operations are performed:
 - a. Create a Layer 2 aggregate interface on the VTEP.
 - b. On the Layer 2 aggregate interface, configure MAC-based access control for 802.1X, and enable MAC-based traffic match mode for dynamic Ethernet service instances.
 - c. Save the configuration and reboot the switch.

201711130686

- Symptom: After the 5930 24-port 10GBASE-T and 2-port QSFP+ with MACsec Module JH182A interface card is rebooted, interfaces on it cannot come up.
- Condition: This symptom occurs if the following operations are performed:
 - a. In an IRF 3.1 system, the 5940 2-slot Switch JH397A/5940 4-slot Switch JH398A switch has an 5930 24-port 10GBASE-T and 2-port QSFP+ with MACsec Module JH182A interface card installed. Configure the interface numbered 4, 8, 12, 16, 20, or 24 as an IRF physical interface.
 - b. Reboot the 5930 24-port 10GBASE-T and 2-port QSFP+ with MACsec Module JH182A interface card.

201710210081

- Symptom: In a VXLAN network, the multicast suppression bandwidth configuration does not take effect on a VSI.
- Condition: This symptom occurs if the following operations are performed:
 - a. On a VSI, configure the **igmp-snooping drop-unknown** command.
 - b. On another VSI, configure the **restrain multicast** command.

201709290520

- Symptom: The outbound packet statistics for a VSI interface are incorrect.
- Condition: This symptom occurs if the VSI interface forwards Layer 3 packets.

201803070942

- Symptom: Packets without VLAN tags cannot be forwarded at Layer 3.
- Condition: This symptom occurs if the following conditions exist:
 - The switch operates in border mode.
 - An interface on the switch is split into four breakout interfaces.
 - On a breakout interface, use the **port trunk permit vlan** command to assign the interface to multiple VLANs as a trunk port.

201803050415

- Symptom: When both the input port and output port of an OpenFlow flow entry are an aggregate interface, the packets received on a member port of the aggregate interface are forwarded out of another member port.

- Condition: This symptom occurs if the following operations are performed:
 - a. Create an aggregation group, and assign multiple ports to the aggregation group.
 - b. Enable the global mode for an OpenFlow instance.
 - c. Activate the OpenFlow instance.
 - d. The controller deploys a flow entry, with both the input port and output port as the aggregate interface.
 - e. A member port of the aggregate interface receives packets.

201802280788

- Symptom: When a 40-GE interface on the switch is configured with cut-through forwarding, packets with CRC checksum errors are dropped.
- Condition: This symptom occurs if cut-through forwarding is enabled on a 40-GE interface of the switch.

201802280161

- Symptom: The NTP service is enabled if the switch configuration is restored by using a .cfg configuration file.
- Condition: This symptom occurs if the following operations are performed:
 - a. Use the **ntp-service unicast-server** command to specify an NTP server for the switch. The **ntp-service enable** command is not configured.
 - b. Save the configuration. Restore the configuration by using the .cfg configuration file.

201802110563

- Symptom: Reading the hh3cBgpEvpn node through MIB failed.
- Condition: This symptom occurs if you read the hh3cBgpEvpn node through MIB.

201802090624/201802030291

- Symptom: Deploying the VRF configuration through NETCONF failed.
- Condition: This symptom occurs if NETCONF is used to deploy the VRF configuration and configure the NexthopVrflIndex field.

201802050047

- Symptom: After a local mirroring group is deleted, the related underlayer entries remain.
- Condition: This symptom occurs if the following operations are performed:
 - a. Create a local mirroring group.
 - b. Configure an aggregate interface as the source port of the local mirroring group.
 - c. Delete the aggregate interface. Then, configure the aggregate interface as the source port of the local mirroring group.
 - d. Delete the mirroring group.

201801250361

- Symptom: When the loop protection action is set to shutdown and the switch detects loops, the **display loopback-detection** command cannot display the related information.
- Condition: This symptom occurs if loop detection is enabled on the switch and the loop protection action is set to shutdown.

201712210251

- Symptom: In a VXLAN network, a VTEP improperly forwards ARP packets carrying double VLAN tags.

- Condition: This symptom occurs if the following operations are performed:
 - a. Enable ARP flood suppression in VSI view.
 - b. The VTEP receives ARP packets carrying double VLAN tags from the AC associated with the VSI.

201711150718

- Symptom: In an IRF 3.1 system, a PEX forward two copies of BUM traffic incorrectly.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure a Layer 2 Ethernet interface as an AC and associate it with a VSI. Assign the interface to a Layer 2 aggregate interface. Then, configure the Layer 2 aggregate interface as an AC and associate it with the same VSI.
 - b. The PEX receives BUM traffic.

201711150550

- Symptom: An IRF fabric splits after the IRF fabric is configured with IGMP snooping and a master/subordinate switchover occurs.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure the IRF fabric with IGMP snooping and a large number of VSIs.
 - b. Perform a master/subordinate switchover for the IRF fabric.

201801120435/201801120438

- Symptom: Failed to configure an aggregate interface or its member port as the source port of a local mirroring group.
- Condition: This symptom occurs if the following operations are performed:
 - a. On an IRF fabric, assign ports numbered the same on different member devices to the same aggregation group.
 - b. Configure the aggregate interface as the source port of a local mirroring group.
 - c. Configure a member port of the aggregate interface as the source port of a local mirroring group.

201802240102/201802240096

- Symptom: The VXLAN tunnels flap.
- Condition: This symptom occurs if the following operations are performed:
 - d. An OVSDB connection is established between an IRF fabric and a controller.
 - e. A master/subordinate switchover occurs to the IRF fabric.

201712070427

- Symptom: The packets sent by the device carry incorrect secondary VLAN tags.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure private VLAN on an IRF fabric.
 - b. Configure the **port private-vlan** command on an aggregate interface.
 - c. The subordinate member device sends packets out of the aggregate interface.

201801120195/201801220358/201801130288/201801220209

- Symptom: Few ARP entries are lost on the switch.
- Condition: This symptom occurs if controllers deploy configuration by using OVSDB to the switch and active/standby switchover is performed for controllers.

201712210968

- Symptom: The BGP neighbors go down and come up.
- Condition: This symptom occurs if a master/subordinate switchover is performed for an IRF fabric when BGP is withdrawing type-5 routes.

201712210740

- Symptom: Layer 3 traffic might fail to be forwarded properly.
- Condition: This symptom occurs if a Layer 3 interface is quickly configured with an IP address and secondary IP addresses.

201712180761/201712180727

- Symptom: In an EVPN network, ARP entries are incorrectly synchronized, and residual MAC address entries exist.
- Condition: This symptom occurs if the following operations are performed:
 - a. Enable ARP flood suppression on a VTEP.
 - b. The VTEP learns ARP information locally.
 - c. The VTEP receives ARP information with the same IP address but a different MAC address.

201801300617/201712110085/201801120501

- Symptom: Directly-connected VPNs bound to subinterfaces with the same subinterface number cannot forward traffic between them.
- Condition: This symptom occurs if the following operations are performed:
 - d. Create Layer 3 subinterfaces R1.2 and R2.2, and bind them to VPNs.
 - e. Remove the binding between subinterface R1.2 and its VPN.

201801241004/201801241000

- Symptom: The memory leaks.
- Condition: This symptom occurs if a VCFC controller is used to deploy configuration to the switch.

201802080387/201802080401

- Symptom: After a VLAN interface is deleted, the Layer 3 Ethernet subinterface or Layer 3 aggregate subinterface with the subinterface number as the VLAN interface number cannot ping a directly-connected device.
- Condition: This symptom occurs if the following operations are performed:
 - a. Create a VLAN and configure a VLAN interface for the VLAN. Then, create a Layer 3 Ethernet subinterface or Layer 3 aggregate subinterface with the subinterface number as the VLAN interface number.
 - b. Delete the VLAN interface.

201802070573

- Symptom: OSPF neighbor relationship cannot be established.
- Condition: This symptom occurs if OSPF is enabled on a VSI interface of the specific network segment.

201802060048

- Symptom: A Layer 3 Ethernet subinterface incorrectly collects the traffic statistics of the corresponding VLAN interface.

- Condition: This symptom occurs if the **traffic-statistic enable** command is not configured on the Layer 3 Ethernet subinterface.

201802010317

- Symptom: The SNMP function fails.
- Condition: This symptom occurs if the following operations are performed:
 - a. Use the **snmp-agent port** command to specify a non-default UDP port for receiving SNMP packets.
 - b. A large number of Layer 3 packets are sent to the CPU.

201801310870

- Symptom: The logs output by the resource monitoring module have errors.
- Condition: This symptom occurs if the total number of MAC addresses on an IRF fabric exceeds the MAC address table size of a single member device, for example, because some traffic is received and sent out only on a certain member device.

201801220376

- Symptom: Traffic cannot be forwarded when the IP address of a multiport ARP entry is the next hop of a route.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure a multiport ARP entry.
 - b. Use the IP address of the multiport ARP entry as the next hop of a route.

201801170954

- Symptom: The **default** command on an interface cannot clear the PFC-related commands on the interface.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure PFC-related commands on the interface.
 - b. Execute the **default** command on the interface.

201801170905

- Symptom: After a 40-GE interface on an LSWM18QC interface card is split, the 10-GE breakout interfaces cannot be displayed.
- Condition: This symptom occurs if the following operations are performed:
 - a. Install a 40GE-to-10GE transceiver module in a 40-GE interface on an LSWM18QC interface card. Execute the **using tengige** command on the interface to split the 40-GE interface into 10-GE breakout interfaces.
 - b. Unplug and plug the interface card, or reboot the interface card.

Resolved problems in R2610

201710310325/201711010685

- Symptom: When the **packet-filter vlan-interface** command is used to apply different types of ACLs to the same direction of two VLAN interfaces for packet filtering, packet filtering does not take effect on the VLAN interface configured later.
- Condition: This symptom might occur if the **packet-filter vlan-interface** command is used to apply different types of ACLs to the same direction of two VLAN interfaces for packet filtering.

201711010631/201710260063/201710260080

- Symptom: On an EVPN network, EVPN packets cannot be forwarded through an Ethernet service instance when MAC address learning is disabled for the Ethernet service instance.
- Condition: This symptom might occur if MAC address learning is disabled for the Ethernet service instance.

201711010692/201710300031

- Symptom: An ACL that matches the inner Ethernet header of VXLAN packets is used for packet filtering. After the rules of the ACL are deleted and reconfigured, packet filtering does not take effect.
- Condition: This symptom might occur if an ACL that matches the inner Ethernet header of VXLAN packets is used for packet filtering, and the rules of the ACL are deleted and reconfigured.

201711060065/201711030386/201711060106

- Symptom: Fragmented UDP packets received from customer sites cannot be transmitted between two VXLAN VTEPs.
- Condition: This symptom might occur if VXLAN VTEPs transmit fragmented UDP packets received from customer sites.

201710270312/201710270307/201711010677

- Symptom: On a distributed EVPN gateway, ARP entries cannot be deleted completely if the gateway receives Layer 3 VXLAN traffic from the local site and ARP requests with the same destination IP address from a remote device in sequence.
- Condition: This symptom might occur if a distributed EVPN gateway receives Layer 3 VXLAN traffic from the local site and ARP requests with the same destination IP address from a remote device in sequence.

201711010633/201710250739

- Symptom: On a VXLAN network, MAC address entries for a VSI cannot be learned after MAC address learning is re-enabled for the VSI.
- Condition: This symptom might occur if the following operations are performed:
 - a. Disable MAC address learning for a VSI.
 - b. Save the running configuration and reboot the device.
 - c. Enable MAC address learning for the VSI.

201711010669/201710240723

- Symptom: On a VXLAN VTEP, port isolation configuration cannot be deleted from a Layer 2 aggregate interface.
- Condition: This symptom might occur if a Layer 2 aggregate interface is assigned to a port isolation group and then is removed from that group on a VXLAN VTEP.

201711060574/201711060601

- Symptom: The switch reboots unexpectedly when a VPN instance of MPLS L3VPN is disassociated from the VLAN interface of a primary VLAN.
- Condition: This symptom might occur if the following operations are performed:
 - a. Associate a VPN instance of MPLS L3VPN with VLAN interface A.
 - b. Configure VLAN A as a primary VLAN and configure VLAN B as a secondary VLAN of VLAN A.
 - c. Disassociate the VPN instance from VLAN A.

201711010640/201710240669

- Symptom: On a VXLAN network, learned MAC address entries for a VSI are not deleted after MAC address learning is disabled for the VSI.
- Condition: This symptom might occur if MAC address learning is enabled and then disabled for the VSI.

201711250125

- Symptom: An IRF fabric acts as a VTEP in an EVPN network. During an ISSU from R2509P02, R2510P01, R2510P02, or E2603 to a later version for the IRF fabric, Layer 3 EVPN traffic loss occurs.
- Condition: This symptom might occur if an ISSU is performed to upgrade an IRF fabric that acts as an EVPN VTEP from R2509P02, R2510P01, R2510P02, or E2603 to a later version.

201711170089/201711170074

- Symptom: On a VXLAN network, VMs cannot ping each other.
- Condition: This symptom might occur if the switch acts as a VTEP and more than six aggregate interfaces are outgoing interfaces of VXLAN tunnels.

201711170122/201711130156/201711140528

- Symptom: Two connected VTEPs cannot ping each other after a VXLAN tunnel is established between them and the VXLAN tunnel interfaces are associated with a VPN instance.
- Condition: This symptom might occur if the following operations are performed:
 - a. Associate the VLAN interfaces on the VTEPs with a VPN instance, and assign an IP address to each VLAN interface.
 - b. Ping a VTEP from the other VTEP by using the **ping** command with the **vpn-instance vpn-instance-name** option specified. The peer VTEP is reachable.
 - c. Establish a VXLAN tunnel between the VTEPs.
 - d. Associate the VXLAN tunnel interfaces on both the VTEPs with the VPN instance, and assign an IP address to each VXLAN tunnel interface.
 - e. Ping a VTEP from the other VTEP by using the **ping** command with the **vpn-instance vpn-instance-name** option specified.

201711010645/201710210020

- Symptom: Users can pass authentication and come online even though the status of the RADIUS server is set to blocked.
- Condition: This symptom might occur if the RADIUS server load sharing feature is enabled.

201711070428/201709250438/201709250432

- Symptom: On a leaf node of a VCF fabric, overlay configuration issued by the VCF controller is lost on a server-facing interface when the cable is re-installed or the interface is shut down and then brought up.
- Condition: This symptom might occur if the VCF fabric is automatically deployed by a VCF controller and VCF Fabric Director.

201711280253/201711290096

- Symptom: In an FCoE network, the switch discards FDISC packets with a sequence count (SEQ_CNT) of 255.
- Condition: This symptom occurs if a node logs in to the switch through a VFC interface and sends FDISC packets to the switch.

201711070855

- Symptom: The **undo jumboframe enable** command does not take effect after an IRF fabric restores a .cfg configuration file.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure the device to prevent jumbo frames from passing through by using the **undo jumboframe enable** command.
 - b. Upgrade the software version of the IRF fabric from R2311P04 to R2422P01 or from R2422P01 to R2432 or later.
 - c. Restore the configuration of the IRF fabric from a .cfg configuration file.

201709280643

- Symptom: The system prompts that memory resources are insufficient.
- Condition: This symptom occurs when the switch is under Telnet attacks.

201711010629/201710180603/201710200121

- Symptom: The connection interfaces on the device and the peer device cannot come up after the interfaces are configured to operate in 1000 Mbps and full duplex mode.
- Condition: This symptom might occur if both interfaces are configured to operate in 1000 Mbps and full duplex mode.

201711170138/201711110207

- Symptom: After an applied QoS policy is removed, the ACL resources used by the QoS policy are not released.
- Condition: This symptom might occur if the following operations are performed:
 - a. Create an aggregation group. No member ports are added to the aggregation group.
 - b. Apply a QoS policy to the aggregation interface.
 - c. Create an ACL and configure rules in the ACL.
 - d. Create a traffic classifier and define a match criterion for the traffic class to match the ACL.
 - e. Modify rules in the ACL.
 - f. Remove the QoS policy applied to the aggregate interface.

201711060893/201711060892

- Symptom: In an EVPN distributed relay system, when a DR interface is shutdown, the VXLAN tunnel that acts as the IPL does not forward data traffic.
- Condition: This symptom might occur if multiple VSIs are configured on the shutdown DR interface.

201710270263

- Symptom: OpenFlow ARP entries cannot be issued when packet loss prevention is enabled for OpenFlow forwarding.
- Condition: This symptom might occur if packet loss prevention is enabled for OpenFlow forwarding.

201710300379

- Symptom: On an IRF 3.1 system, a Layer 2 extended-link aggregation group cannot forward Layer 2 unicast traffic when the aggregation member ports on a PEX are shut down.
- Condition: This symptom might occur if the Layer 2 extended-link aggregation group contains member ports on multiple PEXs.

201711290143

- Symptom: The switch is operating in FCF mode. When an FC interface that is an access port in a VSAN is repeatedly shut down and brought up, FC protocol packets are dropped and users cannot come online on the interface.
- Condition: This symptom might occur if the switch is operating in FCF mode, and the FC interface is an access port in a VSAN.

201710100413/201709220727

- Symptom: On an IRF 3.1 system, high CPU usage of the parent fabric causes LLDP neighbor aging and re-establishment of LLDP neighbor relationships.
- Condition: This symptom might occur if a large number of LLDP neighbors exist.

201710100491/201709250899

- Symptom: If IRF master/subordinate switchovers are performed frequently, the DBM module cannot synchronize data correctly, and the view of interfaces on an IRF member switch cannot be accessed.
- Condition: This symptom might occur if IRF master/subordinate switchovers are performed frequently.

201709010601/201709010596

- Symptom: When the default user role feature is enabled for remote AAA users, the **boot-loader pex file** command fails.
- Condition: This symptom might occur if the default user role feature is enabled for remote AAA users by using **role default-role enable**.

201709040792/201705240732

- Symptom: On an IRF fabric, the console does not respond when configuration is made in interface range view.
- Condition: This symptom might occur if configuration is made in interface range view on an IRF fabric.

201709290115/201709290118

- Symptom: The SNMP process stops working unexpectedly when the **snmp-agent port** command is executed.
- Condition: This symptom might occur if the **snmp-agent port** command is executed.

Resolved problems in R2609

201711070428/201709250438/201709250432

- Symptom: On a leaf node of a VCF fabric, overlay configuration issued by the VCF controller is lost on a server-facing interface when the cable is re-installed or the interface is shut down and then brought up.
- Condition: This symptom might occur if the VCF fabric is automatically deployed by a VCF controller and VCF Fabric Director.

201711060574

- Symptom: The switch reboots unexpectedly when a VPN instance of MPLS L3VPN is disassociated from the VLAN interface of a primary VLAN.
- Condition: This symptom might occur if the following operations are performed:
 - a. Associate a VPN instance of MPLS L3VPN with VLAN interface A.

- b. Configure VLAN A as a primary VLAN and configure VLAN B as a secondary VLAN of VLAN A.
- c. Disassociate the VPN instance from VLAN A.

201711030582

- Symptom: On a distributed-relay member device, the IPL cannot transmit BGP protocol packets if the MTU of the IPP is larger than 2000 bytes.
- Condition: This symptom might occur if the MTU of the IPP is set to a value larger than 2000 bytes on a distributed-relay member device.

201711030386/201711060106

- Symptom: Fragmented UDP packets received from customer sites cannot be transmitted between two VXLAN VTEPs.
- Condition: This symptom might occur if VXLAN VTEPs transmit fragmented UDP packets received from customer sites.

201711020176/201711010332

- Symptom: BGP sessions cannot be established to IPv6 IBGP peers after the **peer ignore** and **undo peer ignore** commands are executed in sequence.
- Condition: This symptom might occur if IPv6 IBGP peers are configured, and the **peer ignore** and **undo peer ignore** commands are executed in sequence.

201711020121

- Symptom: In the output from the **display power** command, the mode of a DC power supply is displayed as HVDC.
- Condition: This symptom might occur if the **display power** command is executed to view information of a DC power module.

201711010881/201710130153

- Symptom: After the **issu one-step** command is executed, the "Do you want to delete flash:/xxx.ipe" message is displayed twice during the ISSU.
- Condition: This symptom might occur if the **issu one-step** command is executed.

201711010734

- Symptom: A Layer 2 extended-link aggregation group with member ports on multiple PEXs is configured on an IRF 3.1 system. After all member ports of the aggregation group are removed and then a member port is re-assigned to the aggregation group, that member port floods received unicast traffic.
- Condition: This symptom might occur if all member ports of a Layer 2 extended-link aggregation group are removed and then a member port is re-assigned to the aggregation group.

201711010700

- Symptom: A VXLAN VTEP broadcasts received remote ARP requests twice to the local site.
- Condition: This symptom might occur if a VXLAN VTEP receives ARP requests on VXLAN tunnel interfaces.

201711010692/201710300031

- Symptom: An ACL that matches the inner Ethernet header of VXLAN packets is used for packet filtering. After the rules of the ACL are deleted and reconfigured, packet filtering does not take effect.
- Condition: This symptom might occur if an ACL that matches the inner Ethernet header of VXLAN packets is used for packet filtering, and the rules of the ACL are deleted and reconfigured.

201711010689

- Symptom: In an EVPN distributed-relay system, one distributed-relay member device has two VXLAN tunnels with the same destination address to the other distributed-relay member device. After the **reset bgp** command is executed on the devices, the IPL cannot recover.
- Condition: This symptom might occur if one distributed-relay member device has two VXLAN tunnels with the same destination address to the other distributed-relay member device in an EVPN distributed-relay system.

201711010685

- Symptom: When the **packet-filter vlan-interface** command is used to apply different types of ACLs to the same direction of two VLAN interfaces for packet filtering, packet filtering does not take effect on the VLAN interface configured later.
- Condition: This symptom might occur if the **packet-filter vlan-interface** command is used to apply different types of ACLs to the same direction of two VLAN interfaces for packet filtering.

201711010677

- Symptom: On a distributed EVPN gateway, ARP entries cannot be deleted completely if the gateway receives Layer 3 VXLAN traffic from the local site and ARP requests with the same destination IP address from a remote device in sequence.
- Condition: This symptom might occur if a distributed EVPN gateway receives Layer 3 VXLAN traffic from the local site and ARP requests with the same destination IP address from a remote device in sequence.

201711010669

- Symptom: On a VXLAN VTEP, port isolation configuration cannot be deleted from a Layer 2 aggregate interface.
- Condition: This symptom might occur if a Layer 2 aggregate interface is assigned to a port isolation group and then is removed from that group on a VXLAN VTEP.

201711010667/201710200410

- Symptom: The site connected to an EVPN distributed-relay system receives two identical ARP packets if ACs are deleted and reconfigured on the Layer 2 aggregate interfaces that act as distributed-relay interfaces.
- Condition: This symptom might occur if the EVPN distributed-relay system uses a VXLAN tunnel as the IPL, and ACs are deleted and reconfigured on the Layer 2 aggregate interfaces that act as distributed-relay interfaces.

201711010664/201710200476

- Symptom: 802.1X authenticated users on a Layer 2 aggregate interface cannot obtain IP addresses through DHCP.
- Condition: This symptom might occur if the Layer 2 aggregate interface is configured with an AC, and the users are assigned the VSI associated with the AC.

201711010663/201710240009

- Symptom: On a VXLAN VTEP, an aggregate interface configured with ACs drops a packet received from an AC when it is both the incoming and outgoing interfaces of that packet.
- Condition: This symptom might occur if bridging is enabled and then disabled on the aggregate interface.

201711010657

- Symptom: On a VPLS network, packets get lost when the packets are forwarded from a PE to a CE.
- Condition: This symptom might occur if the following conditions exist:

- The PE is a multi-chassis IRF fabric.
- The PE receives packets from a peer PE and then a member device in the IRF fabric forwards the packets to a CE.

201711010656/201709160042

- Symptom: A VXLAN tunnel interface goes down if the destination address of the tunnel interface is repeatedly changed.
- Condition: This symptom might occur if the destination address of the tunnel interface is repeatedly changed.

201711010645/201710210020

- Symptom: Users can pass authentication and come online even though the status of the RADIUS server is set to blocked.
- Condition: This symptom might occur if the RADIUS server load sharing feature is enabled.

201711010643/201709130191/201709130181

- Symptom: After an Ethernet service instance is remapped to a VSI, a static MAC address entry for the VSI cannot be recreated because the entry already exists. However, no MAC address entry for the VSI is displayed by using the **display l2vpn mac-address** command.
- Condition: This symptom might occur if the following operations are performed on a VXLAN network:
 - a. Map an Ethernet service instance to a VSI.
 - b. Configure a static MAC address entry for the VSI.
 - c. Remove the mapping between the Ethernet service instance and the VSI.
 - d. Map the Ethernet service instance to the VSI.
 - e. Recreate the static MAC address entry for the VSI.
 - f. Display MAC address entries for the VSI.

201711010642/201709290465

- Symptom: A controller fails to apply a QoS policy to an Ethernet service instance configured to match any frames that do not have an 802.1Q VLAN tag.
- Condition: This symptom might occur if the controller repeatedly applies a QoS policy to an Ethernet service instance configured to match any frames that do not have an 802.1Q VLAN tag.

201711010641/201710180533

- Symptom: On a VXLAN network with a centralized IP gateway, the device acts as the IP gateway. VXLAN packets cannot be forwarded after a primary VXLAN tunnel and a backup VXLAN tunnel are assigned to a VXLAN and the device is rebooted.
- Condition: This symptom might occur after the following operations are performed:
 - a. Assign a primary VXLAN tunnel and a backup VXLAN tunnel are assigned to a VXLAN in VXLAN view.
 - b. Save the running configuration and reboot the device.

201711010640

- Symptom: On a VXLAN network, learned MAC address entries for a VSI are not deleted after MAC address learning is disabled for the VSI.
- Condition: This symptom might occur if MAC address learning is enabled and then disabled for the VSI.

201711010637/201709190854

- Symptom: On a multicast network, the mcsd process on a device configured with PIM snooping might exit unexpectedly.

- Condition: This symptom might occur if the following operations are performed:
 - a. The PIM snooping-configured device forwards multicast packets.
 - b. On a device configured with PIM, clear all multicast forwarding entries and configure OSPF and OSPFv3.
 - c. The PIM snooping-configured device forwards multicast packets for a time.

201711010633/201710250739

- Symptom: On a VXLAN network, MAC address entries for a VSI cannot be learned after MAC address learning is re-enabled for the VSI.
- Condition: This symptom might occur if the following operations are performed:
 - a. Disable MAC address learning for a VSI.
 - b. Save the running configuration and reboot the device.
 - c. Enable MAC address learning for the VSI.

201711010632/201709110803/201708090648/201710130117

- Symptom: Users fail to pass RADIUS authentication and the access process exits unexpectedly when the RADIUS server load sharing feature is enabled.
- Condition: This symptom might occur if the following conditions exist:
 - The RADIUS server load sharing feature is enabled by using the **server-load-sharing enable** command.
 - The system fails to find the IP address of the RADIUS server based on the host name and VPN instance of the server.

201711010631/201710260063/201710260080

- Symptom: On an EVPN network, EVPN packets cannot be forwarded through an Ethernet service instance when MAC address learning is disabled for the Ethernet service instance.
- Condition: This symptom might occur if MAC address learning is disabled for the Ethernet service instance.

201711010629/201710180603/201710200121

- Symptom: The connection interfaces on the device and the peer device cannot come up after the interfaces are configured to operate in 1000 Mbps and full duplex mode.
- Condition: This symptom might occur if both interfaces are configured to operate in 1000 Mbps and full duplex mode.

201711010628/201710100436/201709230345

- Symptom: On a VXLAN network, the device fails to assign user profiles to 802.1X users after 802.1X is re-enabled.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure user profiles.
 - b. Enable 802.1X.
 - c. The device assigns VSIs to online users.
 - d. Disabled 802.1X and re-enable 802.1X.

201711010626/201710200392/201710200404

- Symptom: On an EVPN network, the automatic VCF fabric configuration on an uplink interface of a leaf node is not deleted after the uplink interface is shut down.
- Condition: This symptom might occur if the following operations are performed:
 - a. Change the device name of the spine node.
 - b. Configure the leaf node to use a template file to perform automatic underlay network configuration.
 - c. Reboot the spine node.

- d. Shut down the uplink interface on the leaf node.

201711010620/201710200432

- Symptom: In a VCF fabric, BGP configuration cannot be automatically issued to uplink interfaces on a leaf node after the uplink interfaces are shut down and then brought up.
- Condition: This symptom might occur if the device uses a template file to perform automatic underlay network configuration and uplink interfaces on the leaf node are shut down and then brought up.

201711010431/201709220200

- Symptom: The device runs out of memory if multiple VLAN interfaces are configured to use DHCP or DHCPv6 for address acquisition and then cancel the configuration on the VLAN interfaces repeatedly.
- Condition: This symptom might occur if multiple VLAN interfaces are configured to use DHCP or DHCPv6 for address acquisition and then cancel the configuration on the VLAN interfaces repeatedly.

201711200237/201711170639

- Symptom: On an EVPN network, an IRF fabric acts as a leaf node. When the master device reboots, VXLAN tunnel interfaces on the IRF fabric go down and come up.
- Condition: This symptom might occur if the IRF fabric establishes VXLAN tunnels with other leaf nodes.

201711180092/201711150354

- Symptom: The console does not respond to user input after the configuration is rolled back.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure VPLS and QinQ on the switch.
 - b. Save the running configuration.
 - c. Roll back the configuration.

201711170138

- Symptom: After an applied QoS policy is removed, the ACL resources used by the QoS policy are not released.
- Condition: This symptom might occur if the following operations are performed:
 - a. Create an aggregation group. No member ports are added to the aggregation group.
 - b. Apply a QoS policy to the aggregation interface.
 - c. Create an ACL and configure rules in the ACL.
 - d. Create a traffic classifier and define a match criterion for the traffic class to match the ACL.
 - e. Modify rules in the ACL.
 - f. Remove the QoS policy applied to the aggregate interface.

201711170122/201711130156

- Symptom: Two connected VTEPs cannot ping each other after a VXLAN tunnel is established between them and the VXLAN tunnel interfaces are associated with a VPN instance.
- Condition: This symptom might occur if the following operations are performed:
 - a. Associate the VLAN interfaces on the VTEPs with a VPN instance, and assign an IP address to each VLAN interface.
 - b. Ping a VTEP from the other VTEP by using the **ping** command with the **vpn-instance vpn-instance-name** option specified. The peer VTEP is reachable.
 - c. Establish a VXLAN tunnel between the VTEPs.

- d. Associate the VXLAN tunnel interfaces on both the VTEPs with the VPN instance, and assign an IP address to each VXLAN tunnel interface.
- e. Ping a VTEP from the other VTEP by using the **ping** command with the **vpn-instance vpn-instance-name** option specified.

201711170114/201711160573

- Symptom: The configuration of disabling the device from advertising a prefix does not take effect.
- Condition: This symptom might occur if the **ipv6 nd ra prefix default no-advertise** command is executed to disable the switch from advertising a prefix.

201711170089/201711170074

- Symptom: On a VXLAN network, VMs cannot ping each other.
- Condition: This symptom might occur if the switch acts as a VTEP and more than six aggregate interfaces are outgoing interfaces of VXLAN tunnels.

201711170076

- Symptom: After a storage device registers with an FC switch, the Target information of the storage device in the name service database is incorrect.
- Condition: This symptom might occur if the following conditions exist:
 - a. The FC interface through which the FC switch directly connects to a storage device is shut down and then brought up.
 - b. The storage device sends only one FLOGI request after the interface through which the storage device connects to the FC switch comes up.

201711170075

- Symptom: On a multicast VPN network, deleted multicast forwarding entries remain, and multicast packets cannot be forwarded.
- Condition: This symptom might occur if multiple aggregate interfaces are used as public network interfaces and multiple outgoing interfaces are added to and deleted from multicast forwarding entries in VPN instances.

201711170066

- Symptom: The usage of CPU 0 is nearly 100% when the bC.0 process consumes a large number of CPU resources.
- Condition: This symptom might occur when the bC.0 process consumes a large number of CPU resources.

201711010610/201708180757/201708090760

- Symptom: The BFD configuration cannot be issued to interfaces on an IRF fabric when the IRF fabric performs automatic underlay network deployment.
- Condition: This symptom might occur when the IRF fabric performs automatic underlay network deployment.

201709300398

- Symptom: Packets cannot be forwarded because the next hop in OSPF routes are mistakenly calculated when OSPF neighbors change.
- Condition: This symptom might occur if multiple OSPF neighbors exist in the broadcast domain and the neighbors change.

201711210459

- Symptom: Only the number of packets or the number of bytes is displayed when the traffic statistics for a QoS policy is obtained by using a command or by reading a MIB node.

- Condition: This symptom might occur if a traffic behavior in the QoS policy contains both accounting actions (counting traffic in bytes and packets) and a CAR action.

201711230732

- Symptom: VXLAN traffic fails to be forwarded.
- Condition: This symptom occurs if the link type is configured as trunk for the outgoing interface of a VXLAN tunnel.

201711250125

- Symptom: An IRF fabric acts as a VTEP in an EVPN network. During an ISSU from R2509P02, R2510P01, R2510P02, or F2603 to a later version for the IRF fabric, Layer 3 EVPN traffic loss occurs.
- Condition: This symptom might occur if an ISSU is performed to upgrade an IRF fabric that acts as an EVPN VTEP from R2509P02, R2510P01, R2510P02, or F2603 to a later version.

Resolved problems in R2608P02

201707100377

- Symptom: The console hangs up if the switch rolls back the configuration to a configuration file containing VXLAN and MPLS configuration multiple times.
- Condition: This symptom might occur if the switch rolls back the configuration to a configuration file containing VXLAN and MPLS configuration multiple times.

201706090622

- Symptom: The 5940 2-slot Switch JH397A switch mistakenly generates alarm messages when a 40-GE interface is split into four 10-GE breakout interfaces.
- Condition: This symptom might occur if a 40-GE interface on the 5940 2-slot Switch JH397A switch is split into four 10-GE breakout interfaces.

201707040349

- Symptom: The spanning tree feature fails to discard user traffic on blocked ports, which causes a broadcast storm.
- Condition: This symptom might occur if the blocked ports receive Layer 2 multicast packets.

201706160714

- Symptom: Dynamic MAC-based VLAN assignment does not take effect on a port of the 5940 48XGT 6QSFP28 JH391A switch.
- Condition: This symptom might occur if the spanning tree protocol operates in PVST mode on the switch.

201704220100

- Symptom: A port cannot be assigned to the PVID and the link mode of the port cannot be changed from Layer 2 to Layer 3.
- Condition: This symptom might occur if the following operations are performed:
 - a. Change the link mode of the port between Layer 2 to Layer 3.
 - b. Shut down and then bring up the port.
 - c. Reboot the subcard where the port resides.
 - d. Repeat steps a to c.

201708100372

- Symptom: A 40-GE interface on the 5940 48SFP+ 6QSFP28 JH390A switch cannot come up after the transceiver module on the peer interface is removed and reinserted.

- Condition: This symptom might occur if the following operations are performed:
 - a. Insert a 40-GE transceiver module into the interface. The interface comes up.
 - b. Remove and reinsert the transceiver module on the peer interface.

201708230458

- Symptom: The speed of a 100-GE interface cannot be set to 100000 Mbps or the interface cannot be configured to autonegotiate the speed.
- Condition: This symptom might occur if the following operations are performed:
 - a. Set the speed of the 100-GE interface to 40000 Mbps.
 - b. Insert a 100-GE transceiver module into the interface.

201709050861

- Symptom: MKA sessions cannot be restored after a broadcast storm.
- Condition: This symptom might occur if the following conditions exist:
 - Multiple ports are enabled with MACsec.
 - A broadcast storm occurs on the switch and the MACsec-enabled ports go down and come up repeatedly.

201708190347

- Symptom: In an IRF 3.1 system, when a member port in an aggregation group receives VXLAN packets, the packets are also forwarded out of another member port.
- Condition: This symptom might occur if the aggregate interface is firstly assigned to a port isolation group and then configured with an Ethernet service instance.

201708180643

- Symptom: A 100-GE interface on the 5940 48SFP+ 6QSFP28 JH390A switch cannot come up.
- Condition: This symptom might occur if the interface is connected to a device from a specific vendor and the interface is configured to operate at a specific speed rather than to autonegotiate the speed.

201709010520

- Symptom: In an IRF 3.1 system, traffic cannot be forwarded between PEXs connected to a parent device.
- Condition: This symptom might occur if the cascade port on the parent device is repeatedly shut down and brought up.

201708180556

- Symptom: The system returns error message "Operation failed." after the **loopback external** command is executed.
- Condition: This symptom might occur if the **loopback external** command is executed on an interface to enable external loopback testing.

201709050444

- Symptom: In an IRF 3.1 system, the physical interfaces corresponding to the upstream port on a PEX are mistakenly shut down if the cascade port on the parent device is repeatedly shut down and brought up.
- Condition: This symptom might occur if the cascade port on the parent device is repeatedly shut down and brought up.

201709010116

- Symptom: A port configured with Smart Link cannot send flush messages after MACsec is enabled on it.

- Condition: This symptom might occur after the port is enabled with MACsec.

201708210157

- Symptom: The negotiated rate of an FC interface in down state is displayed as 10G in the output from the **display interface fc [interface-number] brief** command.
- Condition: This symptom might occur if the FC interface is in down state.

201710160757

- Symptom: On a DRNI network, traffic arriving at the IPP of the peer DR member device is also forwarded out of a DR interface of the DR member device.
- Condition: This symptom might occur if the DR member device is rebooted after an Ethernet service instance is configured on the DR interface.

201708280833

- Symptom: After an IRF split, no remote-MAC address entries are displayed in the output from the **display l2vpn mac-address** command though the remote-MAC address entries exist in the driver.
- Condition: This symptom might occur if the following conditions exist:
 - The IRF fabric acts as a VTEP and learns MAC addresses of remote devices.
 - An IRF physical interface is shut down and the IRF fabric splits.
 - The **display system internal overlay mac-address** command is executed in probe view to display remote-MAC address entries in the driver.
 - The **display l2vpn mac-address** command is executed to display MAC address entries for VSIs on the IRF fabric.

201709140924

- Symptom: An IRF fabric cannot restart up if it attempts to load a configuration file containing VPLS configuration.
- Condition: This symptom might occur if the IRF fabric attempts to load a configuration file containing VPLS configuration at startup.

201709110491

- Symptom: After an IRF fabric reboots, IRF physical interfaces come up and then go down and subordinate devices can start up after two reboots.
- Condition: This symptom might occur if the IRF link down report delay is set to 0 milliseconds by using the **irf link-delay 0** command.

201708300506

- Symptom: On a MAN-Core network, the switch acts as a CSR. A memory leakage occurs if a port on an ECMP link is repeatedly shut down and brought up.
- Condition: This symptom might occur if a port on an ECMP link is repeatedly shut down and brought up.

201709010738

- Symptom: In an IRF fabric, an exception occurs in LIPC memory processing, which causes the master and subordinate IRF devices to work improperly.
- Condition: This symptom might occur if the following conditions exist:
 - A process switchover occurs on the IRF fabric.
 - BGP per-prefix label allocation is enabled.

201709300300

- Symptom: On a VXLAN network with distributed IP gateways, a memory leakage occurs on the f6a08ec module.
- Condition: This symptom might occur if the following conditions exist:
 - Multiple aggregate interfaces are used as the traffic outgoing interfaces of a VXLAN tunnel.
 - Member ports of the aggregation groups are shut down and brought up successively.

201609270719

- Symptom: In an IRF fabric, a service port can forward a packet with a length of 10000 bytes but the IRF physical interface cannot.
- Condition: This symptom might occur if a service port receives a packet with a length of 10000 bytes.

201706200741

- Symptom: In an IRF fabric, packet filtering does not take effect on Ethernet service instances.
- Condition: This symptom might occur if packet filtering is configured on Ethernet service instances.

201706290530

- Symptom: URL redirection in portal authentication does not take effect.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure portal authentication.
 - b. Configure the portal Web server to redirect users that request to access URL A to the URL of the portal Web server by using the **if-match** command.
 - c. Configure a portal-free rule to allow HTTP or HTTPS requests destined for the portal Web server to pass through without triggering portal authentication by using the **portal free-rule** command. The portal-free rule takes effect.
 - d. Enter URL A in the address bar of a browser. The Web page is supposed to be redirected to the URL of the portal Web server.

201706170249

- Symptom: An IRF 3.1 system prompts "Operation failed." after the **undo storm-constrain all** command is executed on **a port of a PEX**, but the storm control configuration is removed successfully.
- Condition: This symptom might occur if storm control is configured on the port by using the **storm-constrain command** and then the configuration is removed by using the **undo storm-constrain all** command.

201706160638

- Symptom: The sysname obtained through a RESTCONF API is incorrect.
- Condition: This symptom might occur if a RESTCONF API is used to obtain the sysname.

201707270940

- Symptom: The SecBlade IV ADE card keeps rebooting until the switch hangs up.
- Condition: This symptom might occur if the CPU temperature is too high.

201708170295

- Symptom: The switch reboots unexpectedly after a VSI interface is applied with a QoS policy that contains a traffic accounting action.
- Condition: This symptom might occur if a VSI interface is applied with a QoS policy that contains a traffic accounting action.

201709140947

- Symptom: The filed ip_ecn OpenFlow action set that the controller issues to the switch does not take effect.
- Condition: This symptom might occur if the controller issues the filed ip_ecn OpenFlow action set to the switch.

201708170635

- Symptom: The available outbound traffic policing resources (**EFP meter**) is less than the number of resources supported by the switch.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure a traffic class that uses an ACL as a packet match criterion.
 - b. Configure a traffic behavior that uses a CAR action as a traffic policing action.
 - c. Create a QoS policy and associate the traffic class with the traffic behavior in the QoS policy.
 - d. Apply the QoS policy to outgoing traffic of an interface.
 - e. Configure 129 rules in the ACL. The system prompts insufficient resources.
 - f. Display QoS and ACL resource usage. The output shows that the number of EFP meter resources that has been applied is 128 but the total number of the resources is 512.

201709010769/201709010765

- Symptom: In an IRF 3.1 system, traffic permitted by a QoS policy is duplicated.
- Condition: This symptom might occur if a port of a PEX is applied with a QoS policy that contains the **filter permit** action.

201709040268

- Symptom: An IRF 3.1 system successfully applies an ACL to an interface to filter packets and enables counting ACL rule matches performed in hardware but it prompts the operation fails.
- This symptom might occur if the following operations are performed:
 - a. Apply an ACL to a port on a PEX to filter incoming packets without enabling counting ACL rule matches performed in hardware by using the **packet-filter default inbound** command.
 - b. Apply the same ACL to the same port to filter incoming packets and enable counting ACL rule matches performed in hardware by using **packet-filter default inbound hardware-count** command.

201709070053

- Symptom: An Ethernet interface is applied with a queue scheduling profile that has a queue uses SP, but SP queuing cannot correctly schedule packets on the interface.
- This symptom might occur if the following operations are performed:
 - a. Create a queue scheduling profile by using the **qos qmprofile profile-name** command.
 - b. Configure a queue to use SP by using the **queue queue-id sp** command.
 - c. Repeatedly apply the queue scheduling profile to an Ethernet interface by using the **qos apply qmprofile profile-name** command.

201709080018

- Symptom: The switch reboots unexpectedly after a QoS policy is applied.
- Condition: This symptom might occur if the number of class-behavior associations in the QoS policy exceeds the specification.

201709130740

- Symptom: Layer 2 multicast traffic is flooded in a VLAN even though dropping unknown multicast data is enabled for the VLAN.

- Condition: This symptom might occur if dropping unknown multicast data is enabled in a VLAN and the switch receives Layer 2 multicast traffic for the VLAN.

201708190026

- Symptom: OpenFlow entries are inconsistent with ARP entries and routing entries on a multichassis IRF fabric.
- Condition: This symptom might occur if the IRF fabric is configured with OpenFlow and a master/subordinate is initiated when the controller is issuing or deleting entries.

201708310475

- Symptom: After a master/subordinate switchover, the number of static routing entries in the **display route-static routing-table** command output is inconsistent with the actual number of static routing entries and the excessive entries cannot be deleted.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure MAC-IP flow entries on an aggregate interface.
 - b. Change the outgoing interfaces and destination MAC addresses in all the MAC-IP flow entries.
 - c. Initiate a master/subordinate switchover.

201709050186

- Symptom: Error messages are displayed on the NETCONF client after the client disconnects from the switch.
- Condition: This symptom might occur if the NETCONF client disconnects from the switch.

201708110619

- Symptom: The interface management process on the IRF master device reboots unexpectedly after a master/subordinate switchover.
- Condition: This symptom might occur if the IRF fabric initiates a master/subordinate switchover when receiving IPsec traffic.

201704240287

- Symptom: The snmp process hangs up.
- Condition: This symptom occurs with a low probability if the ifmgr process reboots unexpectedly during a MIB walk.

201709150282

- Symptom: The number of VXLAN tunnels obtained through NETCONF is incorrect.
- Condition: This symptom might occur if the number of VXLAN tunnels is obtained through NETCONF.

201706160122

- Symptom: When the whole IRF fabric is rebooted, the network management interface of the subordinate IRF member device is up, but cannot be pinged.
- Condition: This symptom occurs if the cable is not well connected to the network management interface of the master IRF member device when the whole IRF fabric is rebooted.

201709060253

- Symptom: The service on the physical outgoing interface of a VXLAN tunnel is interrupted, and the south-to-north VXLAN traffic cannot be forwarded.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure VXLAN and enable DHCP snooping.
 - b. The DHCP client sends unicast traffic to the DHCP server.

- c. Shut down the aggregate interface, which is the physical outgoing interface of the VXLAN tunnel.

201709200620

- Symptom: The fiber network management interface cannot come up.
- Condition: This symptom occurs after the device is started.

201708190030

- Symptom: The system prompts that it fails to read information.
- Condition: This symptom occurs if the **display transceiver** command is executed to display transceiver module information about the network management interface.

201709120557

- Symptom: The device is rebooted unexpectedly.
- Condition: This symptom occurs if an aggregate interface acting as the cascade port is deleted from an IRF 3.1 system.

201709260023

- Symptom: The configuration of an aggregate interface on an IRF fabric is lost.
- Condition: This symptom occurs if the following operations are performed:
 - a. Deploy a centralized VXLAN IP gateway network. The IRF fabric acts as a VXLAN IP gateway in the network.
 - b. Perform a master/subordinate switchover on the IRF fabric.

201706290697

- Symptom: The switch is rebooted because the memory is exhausted.
- Condition: This symptom occurs if 100,000 routes exist on the PE and the operations of disabling BGP and rolling back the configuration are repeatedly performed.

201709090216

- Symptom: The switch responds to one ICMP packet with two echo reply packets.
- Condition: This symptom occurs if the following conditions exist:
 - The VXLAN hardware resource mode is set to non-l2gw for the switch.
 - A PBR policy that can redirect ICMP packets to the firewall is issued.

201709040306

- Symptom: In a VXLAN network, Layer 3 traffic cannot be forwarded.
- Condition: This symptom occurs if the following operations are performed:
 - a. Set the VXLAN hardware resource mode to border.
 - b. Configure an interface as an access port, assign it to a VLAN, and configure a MAC address for the VLAN interface of the VLAN.

201708180832

- Symptom: In a VXLAN network, Layer 3 traffic cannot be forwarded in the underlay network.
- Condition: This symptom occurs if a large number of VSI interfaces are repeatedly created and deleted.

201709020270

- Symptom: A policy might not take effect.
- Condition: This symptom occurs if the policy uses an ACL.

201709220231

- Symptom: The **shutdown** command might fail to shut down a port.
- Condition: This symptom occurs if the **shutdown** command is executed to shut down a port when the port is receiving a large number of packets.

201708180844

- Symptom: When SNMP reads information about transceiver modules, the operation times out.
- Condition: This symptom occurs if SNMP is used to read information about transceiver modules of all interfaces of a device with copper ports.

201708170314

- Symptom: In a DRNI network, a multicast packet has two copies.
- Condition: This symptom occurs if DR interfaces are repeatedly assigned to DR groups and removed from DR groups in the DRNI network.

201709250432

- Symptom: The overlay configuration deployed to a port through the VCFC controller is lost.
- Condition: This symptom occurs if the following operations are performed:
 - a. In a VCFC+Director environment, perform auto deployment for the device.
 - b. Unplug and plug the network cable or execute the **shutdown** and **undo shutdown** commands for the interface connecting the leaf switch to the server.

201709040711/201708180749

- Symptom: IRF MAD takes a long time to shut down the upstream port. As a result, OSPF advertises incorrect LSAs.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure static routes, enable LACP MAD, and redistribute OSPF routes on the IRF fabric.
 - b. Perform a master/subordinate switchover.
 - c. The IRF splits.

201708290599

- Symptom: Executing the **undo speed** or **speed 100000** command on a port fails after you install a 40G transceiver module or cable into a 100G port, execute the **speed 40000** command on the port, and remove the transceiver module or cable.
- Condition: This symptom occurs if the following operations are performed:
 - a. Install a 40G transceiver module or cable into a 100G port.
 - b. Execute the **speed 40000** command on the port.
 - c. Remove the transceiver module or cable, and execute the **undo speed** or **speed 100000** command on the port.

201709120655

- Symptom: The minimum guaranteed bandwidth configuration for a queue does not take effect.
- Condition: This symptom occurs if one of the following operations is performed:
 - Configure a QoS queue scheduling profile with the minimum guaranteed bandwidth configuration for a queue and apply it to an outgoing interface.
 - Configure the minimum guaranteed bandwidth for a queue on an interface.

201708280743/201708180717

- Symptom: In an EVPN DRNI network, the AC on one end is configured with static MAC address entries, and the other end correctly issues the synchronized static MAC address entries.

However, the state of these synchronized MAC address entries are displayed as dynamic on the other end.

- Condition: This symptom occurs if the following operations are performed:
 - a. In an EVPN DRNI network, static MAC address entries are configured on the AC at one end.
 - b. Use the **display evpn route mac** command on the other end to display EVPN MAC address entries.

201709060486

- Symptom: When a VXLAN AC receives traffic, the dynamic MAC address entries of the AC are reported as EVPN local MAC address entries. When the dynamic MAC address entries of the AC age out, EVPN local MAC address entries do not age out.
- Condition: This symptom occurs if the following operations are performed:
 - a. The AC receives traffic. The dynamic MAC address entries of the AC are reported as EVPN local MAC address entries.
 - b. Use the **display l2vpn mac-address** command to display MAC address entries.
 - c. Use the **display evpn route mac** command to display EVPN MAC address entries.

201709090106

- Symptom: When a DRNI device shuts down DR interfaces, multiple confirmation messages appear. You need to confirm the operation multiple times.
- Condition: This symptom occurs if the DRNI device shuts down DR interfaces in batches in interface range view.

201708230439

- Symptom: The state of MAC address entries of an aggregate interface on the subordinate IRF member device is incorrect.
- Condition: This symptom occurs if the following operations are performed:
 - a. The aggregate interface on the subordinate IRF member device uses the default access control mode, and it receives Layer 2 traffic.
 - b. Switch the port access control mode of the interface to port-based.
 - c. Use the **display mac-address** command to display the state of MAC address entries.

201704140651

- Symptom: ARP cannot learn entries after certain operations are performed.
- Condition: This symptom occurs if the following operations are performed:
 - a. In MSTP mode, add a Layer 2 Ethernet interface to a Layer 2 aggregation group. The Ethernet interface become Selected.
 - b. After the state of the interface becomes Forwarding, switch the interface to a Layer 3 Ethernet interface and create a Layer 3 Ethernet subinterface on it.
 - c. Switch the spanning tree protocol mode from MSTP to PVST.
 - d. The VLAN to which the Layer 3 Ethernet subinterface belongs is blocked by PVST.

201706220237

- Symptom: Residual routes exist after certain operations are performed.
- Condition: This symptom occurs if the following operations are performed:
 - a. Reboot the whole IRF fabric. Deploy 1k flow entries after the IRF fabric becomes stable.
 - b. Modify the outgoing interfaces and destination MAC addresses of flow entries. Modify the outgoing interfaces to physical interfaces on the master IRF member device.
 - c. Perform a master/subordinate switchover.
 - d. After the IRF fabric is re-formed, view the routing table.

201706120239

- Symptom: The NAS-ID displayed on the device is incorrect after an 802.1X user comes online on a port with port security enabled.
- Condition: This symptom occurs if the following operations are performed:
 - a. Enable port security on an interface, and set the port security mode to **userLoginSecureExt**.
 - b. Execute the **port-security nas-id-profile** command in interface view to specify the NAS-ID profile for the interface.
 - c. In NAS-ID profile view, execute the **nas-id bind vlan** command to bind a NAS-ID to a VLAN.
 - d. A 802.1X user passes authentication and comes online.

201706170332

- Symptom: After an MPLS TE tunnel is deleted through configuration rollback, the subordinate IRF member device reboots.
- Condition: This symptom occurs if the configuration with MPLS TE tunnel configuration is rolled back to the configuration without MPLS TE tunnel configuration.

201707140289

- Symptom: With CRLSP hot backup configured, traffic is not switched back to the primary CRLSP after the primary CRLSP recovers from a failure.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure CRLSP hot backup.
 - b. The primary CRLSP fails. Traffic is switched to the backup CRLSP.
 - c. The primary CRLSP recovers.

201706230542

- Symptom: Traffic is incorrectly forwarded between ACs.
- Condition: This symptom occurs if the following operations are performed:
 - a. In an IRF3.1 system, two PEXs are configured with ACs.
 - b. Save the configuration, and perform a master/subordinate switchover or reboot the whole system.

201706260779

- Symptom: The **arp sender-ip-range** configuration still exists after the configuration is rolled back.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure the **arp sender-ip-range** command for VLAN 1 in the running configuration.
 - b. Roll back the running configuration to the configuration file that does not contain the **arp sender-ip-range** configuration.

201707240706

- Symptom: The system prompts that resources are insufficient when 16K multicast groups are created on the device.
- Condition: This symptom occurs if 16K multicast groups are created on the device.

201705110807

- Symptom: uRPF configuration is lost.
- Condition: This symptom occurs if the following operations are performed:
 - a. Globally enable uRPF on an IRF fabric.
 - b. Perform a master/subordinate switchover.

201708190345

- Symptom: The device might reboot unexpectedly.
- Condition: This symptom occurs if the subcards of the device are rebooted.

201709220383

- Symptom: ACL resources are not released.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure PBR for a VXLAN tunnel.
 - b. Reset the OSPF process used for setting up the VXLAN tunnel.

201708250394

- Symptom: SP queuing on an outgoing interface does not take effect.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure the incoming interface to trust the packet priority.
 - b. Congestion occurs on the outgoing interface.
 - c. Configure the **qos sp** command on the outgoing interface.

201708220134

- Symptom: Residual silent MAC address entries exist.
- Condition: This symptom occurs if the following operations are performed:
 - a. ADCampus is deployed.
 - b. A user fails MAC authentication.
 - c. A user passes 802.1X authentication by using the same MAC address.

201708250117

- Symptom: A ping operation cannot trigger MAC authentication.
- Condition: This symptom occurs if the following operations are performed:
 - a. Enable MAC authentication, and execute the **mac-authentication carry user-ip** command on an interface to include user IP addresses in MAC authentication requests.
 - b. Ping the device from a client.

201709220501

- Symptom: The device might not respond.
- Condition: This symptom occurs if the device has been running for a long period of time.

201709200092

- Symptom: A PEX repeatedly comes online and goes offline.
- Condition: This symptom occurs if the member ports of the upstream port (an aggregate interface) on the PEX go down and come up repeatedly.

201709160042

- Symptom: The state of a VXLAN tunnel becomes down.
- Condition: This symptom occurs if the destination of the VXLAN tunnel is repeatedly modified.

201709080171

- Symptom: A QoS policy cannot be correctly applied.
- Condition: This symptom occurs if the following operations are performed:
 - a. Specify the next startup configuration file for the device, and reboot the device.
 - b. The configuration file contains a lot of traffic accounting and rate limiting configurations. As a result, the QoS and ACL resources are insufficient.

201709150444

- Symptom: MACsec is interrupted after running for a period of time.
- Condition: This symptom occurs if MACsec is enabled on the device and the device forwards a large amount of traffic.

201706100430

- Symptom: Mirroring configuration does not take effect. The destination port cannot receive mirrored packets.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure Layer 2 remote port mirroring (in configurable reflector port method) in the inbound direction.
 - b. The source ports of port mirroring receive packets that need Layer 3 forwarding.

201708190082

- Symptom: A Layer 2 ACL configured with source or destination MAC address rules cannot match unknown unicast traffic.
- Condition: This symptom occurs if a Layer 2 ACL is configured to match unknown unicast traffic.

201709070674

- Symptom: DSCP-based PFC does not take effect on underlay packets.
- Condition: This symptom occurs if the following operations are performed:
 - a. Unicast packets with DSCP 63 and 802.1p 5 are forwarded.
 - b. PFC is enabled for 802.1p values 5 and 7 on the local interface and remote interface.
 - c. Both the local interface and remote interface are configured to trust the DSCP values.

201708280032

- Symptom: Configuration fails.
- Condition: This symptom occurs if the **import** command is used to configure the 802.1p-local priority map.

201708310071

- Symptom: The device reboots unexpectedly after ACL rules are dynamically modified.
- Condition: This symptom occurs if the following operations are performed:
 - a. Use the controller to deploy the packet-filter configuration to a device interface.
 - b. Dynamically modify the ACLs used in the packet-filter.

201708300421

- Symptom: On an IRF fabric, the priority of packets is modified to 0 on the cascade ports. As a result, interface queuing does not take effect.
- Condition: This symptom occurs if packets are received on a parent device and forwarded out of an interface on a PEX.

201708220120

- Symptom: A port does not forward traffic, and the port restores forwarding traffic after the device is rebooted.
- Condition: This symptom occurs if the **cut-through enable** command and then the **undo cut-through enable** command are executed on the device.

201709300037

- Symptom: The PFC feature does not take effect on 5940 devices.

- Condition: This symptom occurs if a port is properly configured and enabled with PFC, and the port receives Layer 3 traffic.

201706210648

- Symptom: An error is returned when you configure a storm suppression threshold higher than the actual bandwidth (negotiated speed) of an interface.
- Condition: This symptom occurs if the following operations are performed:
 - a. The speed of a 100-GE interface is autonegotiated as 40 Gbps or the speed of a 10-GE interface is autonegotiated as 1000 Mbps or 100 Mbps.
 - b. Configure a storm suppression threshold higher than the actual bandwidth (negotiated speed).

201706260671

- Symptom: A Layer 3 Ethernet subinterface cannot ping a directly connected device.
- Condition: This symptom occurs if the directly connected interface of the device is a Layer 3 Ethernet interface, a Layer 3 Ethernet subinterface is created on the interface, and the subinterface is configured with an IP address.

201707030082

- Symptom: After a WRED table configured with ECN is applied to an interface, the ECN function does not take effect.
- Condition: This symptom occurs if a WRED table is created and configured with the ECN function.

201706300189

- Symptom: When an AC configured with the **encapsulation untagged** command and an ACL rule configured with the **permit ip** command are configured on an aggregate interface, the aggregate interface does not forward packets.
- Condition: This symptom occurs if the following operations are performed:
 - a. The controller deploys untagged AC configuration (AC configured with the **encapsulation untagged** command) to the aggregate interface.
 - b. Display the ACL and packet-filter bindings on the aggregate interface.

201709110907

- Symptom: In a VXLAN network, packet loss might occur.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure a VXLAN tunnel on the device.
 - b. The next hop of the tunnel is two aggregation groups.
 - c. Delete one of the aggregation groups.

201709070830

- Symptom: The CLI does not respond when the **follow** command is executed in probe view.
- Condition: This symptom occurs if a process on the device exits exceptionally and the state of process becomes stopped.

201709040776/201705230651

- Symptom: A user fails MAC authentication.
- Condition: This symptom occurs if the following operations are performed:
 - a. Enable port security and set the port security mode to **userlogin-secure-or-mac-ext** on an interface.
 - b. The user fails 802.1X authentication and passes MAC authentication.
 - c. Execute the **shutdown** and **undo shutdown** commands on the interface.

201708280714

- Symptom: After the device is rebooted, part of the mirroring configuration might be lost.
- Condition: This symptom occurs if local port mirroring is configured on an interface, the configuration is saved, and the device is rebooted.

201708250259/201707140139/201708240581

- Symptom: Broadcast storm might appear on the IRF physical interfaces.
- Condition: This symptom occurs if multiple devices form a ring-topology IRF fabric.

201708230635

- Symptom: After the remote MAC address learning function is disabled, MAC address entries will be deleted in the underlayer and MAC address entry information will not be synchronized from the remote end.
- Condition: This symptom occurs if the following operations are performed:
 - a. The MAC address entries are configured not to age out on the device. Packets are normally forwarded.
 - b. Execute the **vxlan tunnel mac-learning disable** command to disable remote MAC address learning.
 - c. Use the **display l2vpn mac-address vsi vsi-name count** command to display MAC address entries for VSIs.

201708180760

- Symptom: In a VXLAN DRNI network, two copies of one broadcast packet are forwarded.
- Condition: This symptom occurs if the following operations are performed:
 - a. A distributed aggregate interface acts as the IPP. Configure an Ethernet service instance with the **encapsulation default** command on the DR interface.
 - b. Save the configuration and reboot the device.

201706260385

- Symptom: When the **undo jumboframe enable size** command is executed on an interface, the system prompts "Operation failed."
- Condition: This symptom occurs if the interface is not configured with jumbo frame support.

201706120372

- Symptom: The memory leaks.
- Condition: This symptom occurs if the **shutdown** and **undo shutdown** commands are executed on the IPP in a DRNI network when packets are being properly forwarded.

201706200784

- Symptom: When the flow entries are repeatedly deployed and deleted, the memory usage reaches the threshold (the memory is released when the free memory reaches 7%).
- Condition: This symptom occurs if the flow entries are repeatedly deployed and deleted.

201708010803

- Symptom: The console port exits and the core file of comsh is generated.
- Condition: This symptom occurs if a large number of commands are frequently issued to the console port.

201708190020

- Symptom: In a DRNI system, EVPN traffic cannot be switched to the new outgoing interface of the IPL after the IPL changes.

- Condition: This symptom might occur if the downlink aggregate interface of the DRNI goes down and the outgoing interface of the IPL changes.

201708300423

- Symptom: In an IRF 3.1 system, the 802.1p priority of packets that a parent device forwards to the connected PEXs is modified to 0.
- Condition: This symptom might occur if the cascade port on the parent device is configured to trust the 802.1p priority.

201708170250

- Symptom: The FCoE process on the switch cannot operate correctly.
- Condition: This symptom might occur if the switch connects to a peer switch through a VFC interface, an FCoE mode is configured on the peer switch, and then the FCoE mode configuration is removed.

201709270412

- Symptom: In a multichassis IRF fabric, ND entries issued by the controller cannot be used to guide packet forwarding after a master/subordinate switchover.
- Condition: This symptom might occur after a master/subordinate switchover.

201708240233

- Symptom: No prompt of IPv6 address conflict is generated when two interfaces are configured to use the same IPv6 prefix to generate the same IPv6 address, and the IPv6 address configured for the second interface does not take effect.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure a static IPv6 prefix by using the **ipv6 prefix** command.
 - b. Configure both interfaces A and B to use the prefix to generate the same IPv6 address by using the **ipv6 address prefix-number** command.

201708210513

- Symptom: The switch reboots unexpectedly.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure an attack defense policy on the switch.
 - b. Launch an ACK flood attack and an ICMP flood attack on the switch.
 - c. Delete the attack defense policy on the switch and clear attack detection and prevention statistics for the switch during the ACK flood attack and ICMP flood attack.

201707110147

- Symptom: In an IRF 3.1 system, when a PEX is rebooted and comes online, the PEX will go offline and then come online.
- Condition: This symptom occurs if a PEX in an IRF3.1 system is rebooted.

201706240115

- Symptom: In an IRF 3.1 system, the configuration of splitting a 40-GE interface into four 10-GE breakout interfaces on a PEX gets lost after the PEX reboots.
- Condition: This symptom might occur after the PEX reboots.

201706120175

- Symptom: In an IRF 3.1 system, the vrdb process exits unexpectedly.
- Condition: This symptom might occur if PEX local forwarding is repeatedly enabled and disabled for a PEX when the PEX keeps receiving known unicast traffic.

201707270657

- Symptom: In an IRF 3.1 system, some member ports of a Layer 2 extended-link aggregate interface on a PEX cannot correctly forward traffic.
- Condition: This symptom might occur if the Layer 2 extended-link aggregate interface has more than one member port.

201709140391

- Symptom: It takes a long for SNMP to return the value of MIB node hh3cTransceiverInfoEntr.
- Condition: This symptom might occur if SNMP is used to read the value of MIB node hh3cTransceiverInfoEntr.

201706220753

- Symptom: Some interfaces on a PEX do not forward traffic.
- Condition: This symptom occurs if VXLAN AC-related operations are repeatedly performed for interfaces on the PEX.

201708020461

- Symptom: Traffic cannot be forwarded.
- Condition: This symptom occurs if the following operations are performed:
 - a. Establish OSPF neighbor between devices.
 - b. Enable MPLS on devices and establish LDP neighbors between devices.
 - c. Configure interfaces connecting these devices to other devices as ACs and an AC receives broadcast traffic.

201706300271

- Symptom: Interfaces flap constantly.
- Condition: This symptom occurs if L3VPN traffic with TTL 1 in the IP header is received.

201704210122

- Symptom: The link mode of an interface cannot be changed between Layer 2 and Layer 3 after MACsec is enabled on the interface.
- Condition: This symptom might occur after MACsec is enabled on the interface.

201706070787

- Symptom: Traffic cannot be forwarded.
- Condition: This symptom occurs if the **shutdown** and **undo shutdown** commands are executed on the aggregate interfaces where the AC and PW reside.

201706140111

- Symptom: Traffic from part of the VLANs cannot be forwarded.
- Condition: This symptom occurs if the device is configured with 4094 VLANs.

201706300173

- Symptom: Loops appear on the parent device in an IRF 3.1 system, and packets are broadcast.
- Condition: This symptom occurs if IGMP is enabled for the same VLAN on the IRF 3.1 system and its neighbor device.

201706300288

- Symptom: Traffic cannot be forwarded by an AC.
- Condition: This symptom occurs if the following operations are performed:
 - Enable PVST on the device.

- The access VLAN of the AC is the same as the S-VLAN ID to be matched by the AC.

201708230476

- Symptom: In an EVPN network, remote MAC address entries cannot be synchronized to the local end by using EVPN routes.
- Condition: This symptom occurs if the following operations are performed:
 - Remote MAC address entries are learned through tunnels.
 - Configure MAC address entries not to age out, and disable remote MAC address learning.

201708190318

- Symptom: The PVID of the public network outgoing interface is different from the VLAN ID carried in the packets forwarded. As a result, the peer PE will drop the received packets.
- Condition: This symptom occurs if the public network outgoing interface is configured as a trunk port in a VPLS network.

201707030578

- Symptom: In a VCF fabric network, a newly added node cannot get the related configurations of an existing router in OpenStack. As a result, the overlay configuration of the node is incorrect.
- Condition: This symptom occurs if the following operations are performed:
 - a. In a VCF fabric network, the device acts as a leaf node and performs automated overlay configuration.
 - b. After the network topology is stable, create a router in OpenStack, and deploy the configuration to nodes with L3 agent enabled in the network.
 - c. Newly add a node, and enable L3 agent on the node.

201708010585

- Symptom: A PEX does not support assigning the MAC learning priority on an interface. However, the corresponding command (**mac-address mac-learning priority high**) can be configured after certain operations are performed.
- Condition: This symptom occurs if the following operations are performed:
 - a. Assign an interface on a PEX to an aggregation group.
 - b. Execute the **mac-address mac-learning priority high** command on the interface assigned to the aggregation group.
 - c. Remove the interface from the aggregation group.

201708190372

- Symptom: The **boot-loader pex** command fails to upgrade the software for PEXs in an IRF 3.1 system.
- Condition: This symptom occurs if the **boot-loader pex** command is executed on a parent device in an IRF 3.1 system to upgrade the software for PEXs.

201709060699

- Symptom: An MPLS TE tunnel cannot forward traffic.
- Condition: This symptom occurs if the following conditions exist:
 - a. Establish an MPLS TE tunnel between two devices by using RSVP-TE, enable the RSVP hello extension function, and configure GR.
 - b. The device receives a large number of L3VPN packets of the specific type, and the routes between the two devices are interrupted.

201708230449

- Symptom: An FC aggregation group member port cannot come up.

- Condition: This symptom occurs if the FC interface is removed from an FC aggregation group and then assigned to the FC aggregation group.

201708180758

- Symptom: The MAC address entries displayed by using the **display mac-address** command are inconsistent with those saved in the device chip.
- Condition: This symptom occurs if the device has learned a large number of MAC address entries and then the **undo mac-address** command is executed to delete MAC address entries.

201708310784

- Symptom: A Layer 3 Ethernet subinterface is blocked unexpectedly.
- Condition: This symptom occurs if the following operations are performed:
 - a. Enable MSTP. Assign a Layer 2 Ethernet interface to a Layer 2 aggregation group. The Ethernet interface becomes Selected.
 - b. After the interface state becomes Forwarding, switch the interface to a Layer 3 Ethernet interface, and create a Layer 3 Ethernet subinterface on it.
 - c. Switch the spanning tree protocol mode from MSTP to PVST.

201709080278

- Symptom: When a DR member device in a DRNI network is rebooted, the service attached to only the other DR member device is interrupted.
- Condition: This symptom occurs if the following operations are performed:
 - a. In a DRNI network, the service is attached to only one DR member device.
 - b. Reboot the other DR member device.

201708300664/201706150387

- Symptom: In an EVPN DRNI network, a VM receives two ARP replies after initiating one ARP request.
- Condition: This symptom occurs if the VM sends an ARP requests to request the address of a silent host attached to two DR member devices.

201706050212

- Symptom: After a PEX comes online, a 40-GE interface on the PEX is split unexpectedly.
- Condition: This symptom occurs if the following operations are performed:
 - a. In an IRF 3.1 system, a PEX uses a non-40-GE interface as the physical interface of the upstream port. The parent device deploys the interface splitting configuration to the 40-GE interfaces on the PEX.
 - b. Change the physical interface of the upstream port on the PEX to a 40-GE interface. The PEX comes online again.

201706300195

- Symptom: The **port up-mode** command configuration is lost on an interface of a PEX.
- Condition: This symptom occurs if the following operations are performed:
 - a. In an IRF 3.1 system, an interface on a PEX is configured with the **port up-mode** command.
 - b. Reboot the PEX.

201708070552/201707100654

- Symptom: The console port of a parent device in an IRF 3.1 system hangs up.
- Condition: This symptom occurs if the following conditions exist:
 - In an IRF 3.1 system, a tier-2 PEX connected to a tier-1 PEX is not correctly connected to the IRF 3.1 system because of cable connection errors.

- The tier-1 PEX receives traffic from the tier-2 PEX.

201707110565

- Symptom: When a Layer 2 extended-link aggregate interface on a PEX acts as an AC, traffic cannot be matched and forwarded by VXLAN.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure VXLAN for an IRF 3.1 system.
 - b. Configure a Layer 2 extended-link aggregate interface of a PEX as an AC.

201707310421

- Symptom: Interfaces on a PEX do not forward traffic.
- Condition: This symptom occurs if the PEX in an IRF 3.1 system is configured with local forwarding.

201708020398

- Symptom: When a standalone interface (non-aggregate interface) on a PEX acts as an AC, MAC address learning operates incorrectly.
- Condition: This symptom occurs if the following operations are performed:
 - a. In an IRF 3.1 system, an interface on a PEX is configured with VXLAN, bound to a VSI, and receives traffic.
 - b. Execute the **display l2vpn mac-address** command.

201708280516

- Symptom: When certain port security modes are used, MAC authentication does not work after a reboot.
- Condition: This symptom might occur if the following operations are performed:
 - a. Enable port security and configure one of the following modes on a port:
 - macAddressWithRadius、
 - macAddressOrUserLoginSecure、
 - macAddressElseUserLoginSecure、
 - macAddressOrUserLoginSecureExt or
 - macAddressElseUserLoginSecureExt;
 - b. Save the running configuration, and delete the binary (.mdb) configuration file;
 - c. Reboot the switch.

201709140917/201709020430

- Symptom: After certain operations, the SSH session stays occupied and is not released.
- Condition: This symptom occurs if the following conditions exist:
 - After you perform authentication through the SSH client, the client does not continue to request services.
 - No operation is performed when you are prompted to modify the password.

201709110550

- Symptom: A server fails to forward storage traffic.
- Condition: This symptom occurs after an IRF fabric connected to the server performs an master/subordinate switchover.

201709050350/201706070058

- Symptom: In a VXLAN network, the MAC address of interface A is incorrectly moved to interface B in the same VLAN.

- Condition: This symptom occurs if the following operations are performed:
 - a. Move the AC on interface A.
 - b. Change the IP address of interface A.
 - c. Assign the original IP address of interface A to interface B.

201709050336/201707040659

- Symptom: In an environment outlined in Appendix E in RFC 2328, OSPF performs incorrect route calculation.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure a static network route.
 - b. Configure the **import-route static** command.
 - c. Configure another static network route.

201709040825/201706120440/201706220025

- Symptom: Port-based 802.1X authentication might fail if the username request timeout timer is set.
- Condition: This symptom might occur if the following operations are performed:
 - a. Enable 802.1X authentication and configure port-based access control on an interface.
 - b. Execute the **dot1x timer tx-period tx-period-value** command in system view.

201708310802

- Symptom: The switch cannot communicate with some devices by sending packets shorter than 64 bytes.
- Condition: This symptom might occur if the switch sends packets shorter than 64 bytes.

201708220184/201708090740

- Symptom: The memory usage is high when multiple Layer 2 aggregation groups receive heavy multicast traffic.
- Condition: This symptom might occur if multiple Layer 2 aggregation groups receive heavy multicast traffic.

201708220161

- Symptom: LACP flaps when the management VLAN setting is configured and then removed on a member port of a dynamic link aggregation group.
- Condition: This symptom might occur if the management VLAN setting is configured and then removed on a member port of a dynamic link aggregation group.

201707240534

- Symptom: Multiple Layer 3 aggregate subinterfaces are assigned to the same VLAN. After one of the Layer 3 aggregate subinterfaces is deleted, the other Layer 3 aggregate subinterfaces cannot forward traffic.
- Condition: This symptom might occur after one of the Layer 3 aggregate subinterfaces is deleted.

201706270107/201708190354

- Symptom: In an IRF 3.1 system, traffic cannot be forwarded out of a Layer 2 extended-link aggregate interface.
- Condition: This symptom might occur if the following conditions exist:
 - The spanning tree feature is enabled on the IRF 3.1 system.
 - The member ports of the Layer 2 extended-link aggregation group reside on different PEXs.
 - A master/subordinate switchover occurs in the parent fabric.

201707120425

- Symptom: A VM cannot be pinged from an IRF fabric that acts as a VTEP on a VXLAN network by using the **emulate-ping vxlan** command.
- Condition: This symptom might occur if the AC that matches the VM's VLAN and the source interface of the VXLAN tunnel reside on different member devices of the IRF fabric.

201707050237

- Symptom: In an IRF 3.1 system, the status of a PEX extended port displayed on the parent device is different from the status of the corresponding physical interface displayed on the PEX.
- Condition: This symptom might occur if the parent device and the PEX are rebooted at the same time.

201707040465/201706280269/201709070779

- Symptom: A NETCONF response returned by the switch contains garbled characters.
- Condition: This symptom occurs with a low probability if a NETCONF get operation is performed on the switch.

201707030686

- Symptom: The switch cannot send ICMP error messages.
- Condition: This symptom might occur if tracer is used for a long time to test the connectivity of the network where the switch resides.

201706290654

- Symptom: On an EVPN network, a subordinate device in an IRF fabric that acts as a VTEP cannot start up after a reboot.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure L2VPN on the IRF fabric.
 - b. Disable L2VPN on the IRF master device and restart the IRF subordinate device at the same time.

201706290591

- Symptom: Two interfaces on the switch both learn the MAC address of the peer device on a TRILL link.
- Condition: This symptom might occur if a TRILL link is established between the switch and the peer device.

201706280553

- Symptom: In an IRF 3.1 system, a PEX fails to forward packets after a reboot.
- Condition: This symptom might occur if the following conditions exist:
 - The IRF 3.1 system is configured with the spanning tree feature.
 - A Layer 2 extended-link aggregation group is created and extended ports on multiple PEXs are added to the aggregation group.
 - One of the PEXs is rebooted.

201706220817

- Symptom: On an MPLS network, the switch displays IPv6 LSP statistics but no IPv6 LSPs exist on the switch.
- Condition: This symptom might occur if a Layer 3 interface is bound to a cross-connect and then the Layer 3 interface is shut down and then brought up.

201706220025

- Symptom: Users fail 802.1X authentication on an interface unexpectedly.

- Condition: This symptom might occur if the 802.1X authentication interface is configured with port-based access control and the username request timeout timer is set by using the **dot1x timer tx-period** command.

201706210717

- Symptom: On a VXLAN network, the switch fails to ping a VM by using the **emulate-ping vxlan** command.
- Condition: This symptom might occur if the length of ICMP echo quests is specified in the **emulate-ping vxlan** command.

201706210696

- Symptom: In an IRF 3.1 system, a member port in a Layer 2 extended-link aggregation group cannot forward packets.
- Condition: This symptom might occur if member ports of the Layer 2 extended-link aggregation group reside on different PEXs and the member port mentioned above is removed from and then added to the aggregation group.

201706200573

- Symptom: In an IRF 3.1 system, Layer 3 traffic cannot be forwarded between a PEX and a parent device.
- Condition: This symptom might occur if the PEX physical interface on the PEX and the connected PEX physical interface on the parent device are assigned to the same VLAN.

201706190568

- Symptom: In an IRF 3.1 system, VXLAN traffic cannot be correctly forwarded.
- Condition: This symptom might occur if a parent device connects to two PEXs and the PEXs are assigned to different PEX groups.

201706160068

- Symptom: Port isolation does not take effect on multiple ports of a multiport unicast MAC address entry.
- Condition: This symptom might occur if the incoming interface and the multiple ports associated with the destination MAC address of frames are assigned to the same port isolation group.

201706140703/201706130226

- Symptom: Broadcast packets get lost when being forwarded through a VXLAN tunnel.
- Condition: This symptom might occur if two VTEPs establish a VXLAN tunnel between two Layer 3 aggregate interfaces and member ports in the aggregation group are repeatedly shut down and then brought up.

201706140490

- Symptom: In an IRF 3.1 system, known Layer 2 unicast traffic cannot be forwarded by a PEX after the PEX goes offline even though persistent forwarding is enabled.
- Condition: This symptom might occur if PEX local forwarding and PEX persistent forwarding are enabled on the PEX and the cascade port for the PEX is shut down to log off the PEX.

201706100377

- Symptom: On a VXLAN network with a centralized IP gateway, packets get lost when being forwarded between VXLAN tunnels on different subnets.
- Condition: This symptom might occur if packets are forwarded between VXLAN tunnels on different subnets.

201706160376

- Symptom: After certain operations, the **display igmp-snooping** command displays information for a non-existent VSI.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure IGMP snooping on VSI A.
 - b. Remove IGMP snooping configuration from VSI A.
 - c. Create VSI B and then delete it.
 - d. Reconfigure IGMP snooping on VSI A.

201702080285

- Symptom: An ACL is applied to both the inbound and outbound directions of an object for packet filtering. When the ACL rule is changed to one that is not supported by one direction, ACL application for the other direction fails.
- Condition: This symptom might occur if the following operations are performed:
 - a. Apply an ACL to both the inbound and outbound directions of an object for packet filtering.
 - b. Modify the ACL rule.

201706270499

- Symptom: The **display qos-acl resource** command displays incorrect information when an ACL is applied to the outbound direction of an interface for packet filtering.
- Condition: This symptom might occur if an ACL is applied to the outbound direction of an interface for packet filtering.

201706200605

- Symptom: No prompt is displayed when ACL rule configuration fails for ACL resource exhaustion.
- Condition: This symptom might occur if ACLs are configured when ACL resources are exhausted.

201707250556

- Symptom: WRED still takes effect on an interface after the WRED table is deleted from the interface.
- Condition: This symptom might occur if a WRED table is applied to an interface, and ECN is enabled for queues.

201707120598

- Symptom: WRR or WFQ queuing configuration fails on an interface.
- Condition: This symptom might occur if WRR or WFQ queuing is configured on an interface.

201707140117

- Symptom: The switch does not support redirecting traffic to a tunnel interface by using a QoS policy.
- Condition: This symptom might occur if a QoS policy is configured to redirect traffic to a tunnel interface.

201706150799

- Symptom: TRILL packets cannot be forwarded transparently when TRILL is globally disabled.
- Condition: This symptom might occur if the following operations are performed:
 - a. Enable TRILL globally and on interfaces.
 - b. Disable TRILL globally.

201706260765/201707120724

- Symptom: When an ACL that does not contain rules is configured for a portal preauthentication domain, service packets are falsely permitted.
- Condition: This symptom might occur if an ACL that does not contain rules is configured for a portal preauthentication domain.

201707080345

- Symptom: An IRF fabric fails to forward traffic of a VXLAN AC to a VXLAN tunnel.
- Condition: This symptom might occur if the virtual nexthops of a VXLAN tunnel are on multiple IRF member devices, and the virtual nexthops change.

201706270792

- Symptom: The switch reboots unexpectedly if ARP entries are added or deleted when the memory usage reaches the upper limit.
- Condition: This symptom might occur if ARP entries are added or deleted when the memory usage reaches the upper limit.

201706200856

- Symptom: ARP entries cannot be issued to an IRF fabric through OpenFlow.
- Condition: This symptom might occur if the following operations are performed:
 - a. Issue MAC-IP flow entries to an aggregate interface on an IRF fabric.
 - b. Delete the aggregate interface.
 - c. Perform a master/subordinate switchover.
 - d. Re-create the aggregate interface.

201706140655

- Symptom: Configuration of PEX local forwarding and PEX persistent forwarding might fail on an IRF 3.1 system.
- Condition: This symptom might occur if the following conditions exist:
 - a. A PEX goes offline after PEX local forwarding is enabled for it.
 - b. PEX persistent forwarding is enabled when the PEX is getting online.

201706120174

- Symptom: The resources occupied by table-miss flow entries of an OpenFlow instance are not released after the OpenFlow instance is deactivated.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure the default action of table-miss flow entries to forward packets to the normal pipeline for an OpenFlow instance.
 - b. Issue table-miss flow entries.
 - c. Execute **undo active instance** to deactivate the OpenFlow instance.

201706270180

- Symptom: The switch reboots unexpectedly when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - The controller is enabled to receive ARP packets from the switch.
 - L2VPN is enabled.
 - No VSI interfaces exist.

201706150748

- Symptom: On an IRF 3.1 system, a TRILL access port on a PEX cannot establish neighbor relationships.
- Condition: This symptom might occur if TRILL is configured on an IRF 3.1 system, and a TRILL access port is configured on a PEX.

201705260524

- Symptom: Portal users who have gone offline cannot come online again.
- Condition: This symptom might occur if the following conditions exist:
 - Portal roaming is enabled.
 - When portal users are online, the Rule ARP or ND entry feature is disabled for portal clients.
 - Rule ARP entries are not deleted completely after the portal users go offline.

201706120496

- Symptom: When port security is enabled, some users who fail authorization cannot go offline on an interface.
- Condition: This symptom might occur if the following operations are performed:
 - a. Enable port security and set the port security mode to **userlogin-secure-ext** on an interface.
 - b. Execute **port-security authorization-fail offline** in system view.
 - c. Execute **dot1x re-authenticate manual** on the interface in step a.

201706190108

- Symptom: Memory leaks occur if NQA has been running for a period of time.
- Condition: This symptom might occur if NQA has been running for a period of time.

201706130641

- Symptom: An interface that hosts ACs cannot forward traffic after IRF port binding is configured and then removed for it.
- Condition: This symptom might occur if IRF port binding is configured and then removed for an interface that hosts ACs.

201707190413

- Symptom: One DR member device of a distributed-relay system cannot ping a directly connected device.
- Condition: This symptom might occur if a distributed-relay system is configured.

201706230111

- Symptom: On an IRF 3.1 system, multicast traffic cannot be forwarded to the interfaces that host ACs on PEXs.
- Condition: This symptom might occur if IGMP snooping is enabled on VSIs, and outgoing interfaces of multicast traffic are on PEXs.

201706060803

- Symptom: In a VPLS network, the switch adds incorrect VLAN tags to traffic forwarded out of ACs on an interface, which causes forwarding failure.
- Condition: This symptom might occur if two or all of the following criteria are configured for ACs on an interface, and then one of the configured criteria is removed:
 - **default**
 - **untagged**

- **svid** (the matching outer VLAN ID is same as the PVID.)

201706120355

- Symptom: In a VPLS network, when the switch prompts for insufficiency of L2VPN hardware resources, VPLS can work on the control plane, but traffic forwarding fails.
- Condition: This symptom might occur if the control word feature is enabled, and PWs are repeatedly created and deleted.

201707030662

- Symptom: On an IRF 3.1 system, interfaces on a PEX cannot be selected in an aggregation group.
- Condition: This symptom might occur if the parent IRF fabric splits and then reunites.

201706050232

- Symptom: On a leaf node of the VCF fabric 2.5 solution, a security policy does not take effect on an interface not configured with an AC after a VM migrates to the interface from an AC that uses the untagged criterion.
- Condition: This symptom might occur if the leaf node is a 5940 switch, and a VM migrates from an AC that uses the untagged criterion to an interface not configured with an AC.

201706240414

- Symptom: When certain conditions exist, IGMP snooping forwarding entries cannot be deleted completely.
- Condition: This symptom might occur if the following conditions exist:
 - IGMP snooping is configured for a VSI on an IRF fabric, and the ACs of the VSI are on two interfaces of different IRF member devices.
 - The interfaces in step a receive IGMP reports.
 - The VSI is shut down.

201706130453

- Symptom: Forwarding fails for traffic of a MAC-based VLAN. After MAC-based VLAN is disabled and all MAC-to-VLAN entries are deleted, traffic forwarding still fails.
- Condition: This symptom might occur if the following operations are performed:
 - Enable dynamic MAC-based VLAN assignment on ports.
 - Create MAC-to-VLAN entries with an all-F mask.
 - Create MAC-to-VLAN entries with a non-all-F mask and then re-enabled.

201706300324

- Symptom: Installation or reboot of an interface card causes BFD session flapping.
- Condition: This symptom might occur if an interface card is installed or rebooted.

201706100298

- Symptom: An IRF fabric that acts as a centralized VXLAN IP gateway cannot forward Layer 3 traffic after a reboot.
- Condition: This symptom might occur if an IRF fabric that acts as a centralized VXLAN IP gateway is rebooted.

201706290280

- Symptom: A distributed-relay system has traffic loss when certain conditions exist.
- Condition: This symptom might occur if one of the following conditions exists:
 - Member ports of a DR interface go down and come up.

- A DR member device reboots.

201706240074

- Symptom: On an IRF fabric, when static IPv4SG or IPv6SG bindings are configured on a Layer 3 aggregate interface, the master displays the "Failed to add a static binding entry on slot 2. Reason: Not support." message, and subordinates display the "Failed to add a static binding entry on slot 1. Reason: Unknown error." message.
- Condition: This symptom might occur if static IPv4SG or IPv6SG bindings are configured on a Layer 3 aggregate interface.

201707110286

- Symptom: On an IRF 3.1 system, not all VLAN port assignment settings configured on the parent IRF fabric are issued to ports of PEXs. As a result, the parent IRF fabric and PEXs have inconsistent VLAN configuration.
- Condition: This symptom might occur if VLAN port assignment settings are configured on an IRF 3.1 system.

201706200624

- Symptom: An unauthenticated 802.1X user cannot join the guest VLAN.
- Condition: This symptom might occur if the following conditions exist:
 - a. A guest VLAN is configured for 802.1X authentication.
 - b. The **dot1x max-user** command is executed on an interface, and the number of users on the interface reaches the set limit.
 - c. 802.1X users access webpages on network segments that do not require authentication.

201707010176

- Symptom: Configured ACLs are falsely deleted after certain operations are performed.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure a PBR policy and reference ACL 3000.
 - b. Apply the PBR policy to an interface.
 - c. Configure the **rule 1 deny ip** and **rule 2 permit ip** rules in sequence for ACL 3000.

201707180492

- Symptom: A PBR policy does not take effect on a super VLAN interface.
- Condition: This symptom might occur if a PBR policy is applied to a super VLAN interface.

201704240579

- Symptom: The CLI responds slowly when certain operations are performed.
- Condition: This symptom might occur if the following operations are performed:
 - a. Establish a VPLS connection to another device.
 - b. Configure 8 K PWs.
 - c. Use an aggregate interface as the transport-facing interface, and shut down and bring up the aggregate interface.

Resolved problems in F2608

201706090491

- Symptom: Configuration of the secure MAC address limits fails if more than 126 secure MAC address limits are configured on an interface.

- Condition: This symptom might occur if the **port-security max-mac-count** *max-count* [**vlan** [*vlan-id-list*]] command sets more than 126 secure MAC address limits on an interface.

201706140365/201706140156

- Symptom: A 5940 switch cannot learn the ARP information for a server connected to the switch through an aggregate interface.
- Condition: This symptom might occur if the following conditions exist:
 - The 5940 switch is one of the following modules:
 - 5940 32QSFP+ JH396A.
 - 5940 2-slot Switch JH397A.
 - 5940 4-slot Switch JH398A.
 - 5940 48XGT 6QSFP+ JH394A.
 - 5940 48SFP+ 6QSFP+ JH395A.
 - 5940 48XGT 6QSFP28 JH391A.
 - 5940 48SFP+ 6QSFP28 JH390A.
 - The server and the switch are connected by an aggregate link.
 - The server sends ARP packets tagged with VLAN ID 0.

201705270697

- Symptom: Member ports in an aggregation group on a CE cannot become Selected.
- Condition: This symptom occurs if the following operations are performed:
 - a. In a VPLS network, transparent transmission of LACP packets is enabled on PEs to implement link aggregation for CEs cross PEs.
 - b. On PEs, execute the **control-word enable** command in PW class view.
 - c. Execute the **shutdown** command to shut down any aggregation group member port on a CE, or shut down an interface between a P device and PE device without interrupting the network connectivity.
 - d. Execute the **undo shutdown** command to bring up the interface shut down in step c.

201705250545/201705240318

- Symptom: Loops occur when VMs migrate in an EVPN network.
- Condition: This symptom might occur if the following conditions exist:
 - a. 300 K VMs send ARP requests to the connected leaf node and the border device.
 - b. The VMs migrate from the leaf node to the border device.

201703070287

- Symptom: When certain conditions exist, the Telnet connection to the switch is interrupted for tens of seconds.
- Condition: This symptom might occur if the following conditions exist:
 - a. The switch is managed through Telnet.
 - b. An interface that hosts ACs receives a large amount of ARP traffic.
 - c. The **reset arp interface vsi** command is executed in user view.

201705310087

- Symptom: A dynamically learned local ARP entry cannot overwrite a remote ARP entry issued through OVSDB.
- Condition: This symptom might occur if the switch has a remote ARP entry issued through OVSDB and a dynamically learned local ARP entry for the same MAC address.

201703270703

- Symptom: The switch is in an EVPN network that contains two route reflectors. When the switch receives a MAC advertisement route with the label 2 attribute, assertions are output.
- Condition: This symptom might occur if the switch is in an EVPN network that contains two route reflectors, and it receives a MAC advertisement route with the label 2 attribute.

201705230358

- Symptom: The switch fails to forward multicast traffic if the egress interface is an aggregate interface, and the member ports of the aggregate interface change.
- Condition: This symptom might occur if the egress interface of multicast traffic is an aggregate interface, and the member ports of the aggregate interface change.

201705160799

- Symptom: A DR system cannot forward upstream traffic when the number of ECMP route next hops changes from two to one.
- Condition: This symptom might occur if the number of ECMP route next hops changes from two to one.

201705250318

- Symptom: On an IRF fabric, an FC aggregate interface cannot come up when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - The subordinate switch in the IRF fabric is installed with FC interface modules.
 - The member interfaces of the FC aggregate interface are on the subordinate switch.

201705240650

- Symptom: An FC interface module reboots slowly if all its FC interfaces are switched from Ethernet interfaces.
- Condition: This symptom might occur if all FC interfaces of an FC interface module are switched from Ethernet interfaces.

201705170267

- Symptom: An HPE device cannot capture packets after the packet capture package is installed.
- Condition: This symptom might occur if the packet capture package is installed on an HPE device.

201705110698

- Symptom: When link-aggregation traffic redirection is enabled on a DR system, slight traffic loss occurs when traffic is switched from one DR interface to another DR interface.
- Condition: This symptom might occur if the following conditions exist:
 - a. Link-aggregation traffic redirection is enabled on a DR system.
 - b. The DR system receives unicast traffic on DR interfaces.
 - c. A DR member device is rebooted, or a DR interface is shut down and brought up.

201703280313

- Symptom: An IRF fabric enabled with EVPN distributed relay fails to forward traffic when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - The EVPN network has two border devices.
 - Two spine nodes provide access to the external network.

- The EVPN network contains an IRF fabric that is formed by two leaf nodes, and EVPN distributed relay is enabled on the IRF fabric.
- Four ECMP routes exist between the IRF fabric and the spine nodes. When some of routes are being used for traffic forwarding, the traffic egress ports on the IRF fabric are shut down.

201704250437

- Symptom: On an IRF fabric, some static ARP entries are lost after an IRF master/subordinate switchover.
- Condition: This symptom might occur if the following conditions exist:
 - An IRF fabric acts as a VTEP.
 - Each IRF member switch has an interface that hosts Ethernet service instances and 512 static ARP entries for the Ethernet service instances.
 - The .mdb file is deleted after the running configuration is saved.
 - An IRF master/subordinate switchover occurs.

201705030056

- Symptom: The switch cannot correctly forward Layer 2 VXLAN traffic when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - a. The controller issues an AC that matches untagged frames to the switch.
 - b. The AC receives untagged frames.

201705260052

- Symptom: A controller fails to issue Ethernet service instance configurations to the interface that VMs migrate to.
- Condition: This symptom might occur if the following conditions exist:
 - The switch is a leaf node, and VMs are connected to the switch through an interface that hosts an Ethernet service instance that matches untagged frames.
 - VMs migrate to an interface where no Ethernet service instance exists.

201701160321

- Symptom: On an IRF 3.1 system, the link aggregation process becomes abnormal when a certain operation is performed.
- Condition: This symptom might occur if one of the following operations is performed on the IRF 3.1 system:
 - Remove a Layer 2 extended-link aggregate interface from tier-2 PEXs.
 - Reboot an IRF member switch that hosts member ports of a multidevice aggregation group.

201706150750

- Symptom: An IRF fabric that acts as a VXLAN VTEP has a large number of ARP entries. When the master switch is rebooted, traffic is interrupted for several minutes.
- Condition: This symptom might occur if an IRF fabric that acts as a VXLAN VTEP has a large number of ARP entries, and the master switch is rebooted.

201705240040

- Symptom: An IRF fabric fails to forward multicast traffic that is received on multiple IRF members when the number of IGMP forwarding entries exceeds 8 K.
- Condition: This symptom might occur if the number of IGMP forwarding entries exceeds 8 K on an IRF fabric.

201705160193

- Symptom: An IRF fabric configured with EVPN fails to forward Layer 3 traffic when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - An IRF fabric is at the leaf layer, and it has a large number of ARP flood suppression entries.
 - A VCF controller performs configuration synchronization for the IRF fabric when a master/subordinate switchover occurs.

201612210596

- Symptom: After certain operations are performed, AC creation fails on an Ethernet interface.
- Condition: This symptom might occur if the following operations are performed:
 - a. Create a Layer 2 aggregate interface and create about 4 K ACs on the interface.
 - b. Delete the aggregate interface.
 - c. Configure ACs on a former member port of the deleted aggregate interface.

201705220641

- Symptom: Two VTEPs enabled with EVPN distributed relay fail to forward traffic sent from a remote site to the local site when the site-facing interface on a VTEP goes down.
- Condition: This symptom might occur if the following conditions exist:
 - EVPN distributed relay is enabled on two VTEPs.
 - The site-facing interface on a VTEP goes down.

201706030375

- Symptom: When an ISSU is performed, Layer 2 known unicast traffic is flooded by the switch.
- Condition: This symptom might occur if the following conditions exist:
 - The switch has a large number of unicast MAC address entries.
 - An ISSU is performed during Layer 2 traffic forwarding.

201705120037/201507280150/201510160228

- Symptom: A VCF controller fails to issue a VSI to the switch when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - a. When VXLAN preprovisioning is disabled for the switch on the VCF controller, a VLAN-to-VXLAN mapping is issued to the switch through the VCF controller.
 - b. A VM comes online.
 - c. Modify the VLAN-to-VXLAN mapping on the VCF controller.

201705040567

- Symptom: A VPLS CE is connected to a PE through an aggregate link, and the PE is an IRF fabric. When a master/subordinate switchover occurs, the aggregate interface on the CE flaps constantly.
- Condition: This symptom might occur if the following conditions exist:
 - A VPLS CE is connected to a PE through an aggregate link.
 - The PE is an IRF fabric, and a master/subordinate switchover occurs.

201705260619

- Symptom: After certain operations are performed on a 40-GE interface, the switch displays information for the transceiver module that has been installed in and removed from the interface before the current transceiver module.

- Condition: This symptom might occur if the following operations are performed on a 40-GE interface:
 - a. Install a transceiver module and remove it.
 - b. Install a transceiver module of a model different from the one in step a.
 - c. Split the interface into four breakout interfaces.

201704200405

- Symptom: If the **speed 1000** and **undo speed** commands are repeatedly executed on an interface, the interface cannot forward traffic.
- Condition: This symptom might occur if the **speed 1000** and **undo speed** commands are repeatedly executed on an interface.

201703300379

- Symptom: If an interface has been assigned to and removed from an aggregation group, an error might occur when the interface is re-assigned to the aggregation group.
- Condition: This symptom might occur if an interface is assigned to and removed from an aggregation group, and then is re-assigned to the aggregation group.

201610210114

- Symptom: IGMP packets are reported to the controller.
- Condition: This symptom occurs if the controller does not issue flow entries for IGMP packets.

201705100630

- Symptom: In a centralized VXLAN IP gateway group network, if the switch acts as an access layer VTEP, the tunnels automatically established between the switch and the gateway group member VTEPs cannot come up.
- Condition: This symptom occurs if the following operations are performed:
 - a. In an IRF fabric network, the master member is first started, and the **vtep group group-ip member remote member-ip<1-8>** command is executed on the master device to specify a VXLAN IP gateway group and its members. In this case, the VXLAN tunnel to the gateway group is not created.
 - b. Another device is started and joins the IRF fabric as a subordinate member.
 - c. After the subordinate member joins the IRF fabric, a VXLAN tunnel to the gateway group is created on the IRF fabric.
 - d. An IRF master/subordinate switchover occurs.

201705100582/201704260154

- Symptom: A 5940 switch labeled with one of the following product codes might reboot unexpectedly in certain conditions:
 - 5940 32QSFP+ JH396A.
 - 5940 2-slot Switch JH397A.
 - 5940 4-slot Switch JH398A.
 - 5940 48XGT 6QSFP+ JH394A.
 - 5940 48SFP+ 6QSFP+ JH395A.
 - 5940 48XGT 6QSFP28 JH391A.
 - 5940 48SFP+ 6QSFP28 JH390A.
- Condition: This symptom might occur if the following conditions exist.
 - The **speed** command is executed in interface view.
 - The switch is in an overlay network, or the **display counter** command is executed.

201703220242

- Symptom: If the **reset bgp all** command is executed when a large number of BGP-VPN instances exist, BGP instance view cannot be entered, and the system displays **BGPM process is busy. Please try it later..**
- Condition: This symptom might occur if a large number of BGP-VPN instances are configured, and the **reset bgp all** command is executed.

201703170609

- Symptom: In the output from the **display distributed-relay role** command, a DR system contains two primary DR member devices.
- Condition: This symptom might occur if the following operations are performed:
 - a. Set up a DR system.
 - b. Use the **distributed-relay role priority** *priority-value* command to modify the DR role priorities of DR member devices.
 - c. Shut down and bring up the IPP on one DR member device.

201703010246/201703160302

- Symptom: The MAD status might be faulty on an IRF fabric that is deployed through automatic configuration.
- Condition: This symptom might occur if the following conditions exist:
 - a. Automatic configuration is used to set up an IRF fabric and to issue MAD configuration on two switches that start without configuration.
 - b. The switches have the same network node role (leaf or spine) and are connected by multiple links.
 - c. An IRF physical interface is disassociated from an IRF port.

201703090513/201703160301

- Symptom: The switch is deployed through a VCF controller. When a QoS policy that collects VXLAN packet statistics is applied globally, the switch prompts for resource insufficiency and operation failure.
- Condition: This symptom might occur if a QoS policy that collects VXLAN packet statistics is applied globally.

201703170289

- Symptom: A multicast forwarding entry does not contain detailed information about the incoming interface that is an aggregate interface.
- Condition: This symptom might occur if the following conditions exist:
 - A VLAN interface is the incoming interface in a multicast forwarding entry.
 - An aggregate interface is assigned to the corresponding VLAN.

201703100784

- Symptom: If the switch is connected to an upstream ZTE device in an MPLS TE network, RSVP cannot establish an LSP with the ZTE device, and the tunnel to the ZTE device cannot come up.
- Condition: This symptom might occur if the switch is connected to an upstream ZTE device in an MPLS TE network.

201703100656

- Symptom: After certain operations are performed, an IRF fabric cannot ping a directly connected peer through a VLAN interface.
- Condition: This symptom might occur if the following operations are performed on an IRF fabric:
 - a. Use the **irf mac-address** *mac-address* command to set the IRF bridge MAC address.

- b. Assign a MAC address to the VLAN interface.
- c. Use the **undo mac-address** command to remove the MAC address from the VLAN interface.
- d. Execute the **undo irf mac-address** command.

201703080528

- Symptom: After certain operations, EVPN cannot forward some Layer 3 packets.
- Condition: This symptom might occur if the following operations are performed:
 - a. Remove the VSI interface associated with an L3 VXLAN ID.
 - b. Execute the **reset bgp all** command.
 - c. Re-create the VSI interface and associate it with the original L3 VXLAN ID.

201703070015

- Symptom: A VCF fabric has two border devices. If all VPN instances use the same RD, tunnels to the border devices cannot be established.
- Condition: This symptom might occur if a VCF fabric has two border devices, and all VPN instances use the same RD.

201702280272

- Symptom: In the output from the **display device manuinfo** command, information about 650 W and 300 W power supplies is incomplete.
- Condition: This symptom might occur if the switch uses 650 W and 300 W power supplies, and the **display device manuinfo** command is executed.

201610210366

- Symptom: On a VXLAN border device, the ACL applied to a VLAN interface does not take effect if the physical outgoing interface of a VXLAN tunnel is in the corresponding VLAN.
- Condition: This symptom might occur if the following operations are performed on a VXLAN border device:
 - a. Assign the physical outgoing interface of a VXLAN tunnel to a VLAN.
 - b. Apply an ACL to the corresponding VLAN interface to deny all IP packets to pass.

201702210178

- Symptom: When the management Ethernet interface is shut down and then brought up, the switch is disconnected from the OpenFlow controller twice.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure OpenFlow.
 - b. Execute the **shutdown** and **undo shutdown** commands in sequence on the management Ethernet interface.

201701160273

- Symptom: The switch cannot learn ARP entries on an Ethernet service instance if the interface that hosts the Ethernet service instance does not permit the outer VLAN that matches the Ethernet service instance.
- Condition: This symptom might occur if the following conditions exist:
 - Both VXLAN and PVST are configured on the switch.
 - The interface that hosts an Ethernet service instance does not permit the outer VLAN that matches the Ethernet service instance.

201702080158

- Symptom: The switch is a spine node in a VCF fabric. When the **display vcf-fabric underlay autoconfigure** command is executed after a large number of leaf nodes go offline, the CLI might stop responding.
- Condition: This symptom might occur if the following conditions exist:
 - The switch is a spine node in a VCF fabric.
 - The **display vcf-fabric underlay autoconfigure** command is executed after a large number of leaf nodes go offline.

201703060165

- Symptom: A long delay occurs when the management Ethernet interface is pinged.
- Condition: This symptom might occur if the management Ethernet interface is pinged.

201703200244

- Symptom: In an EVPN network, some upstream packets sent by leaf nodes contain an incorrect destination MAC address, which causes forwarding failure.
- Condition: This symptom might occur if the following conditions exist:
 - Leaf nodes are connected to spine nodes through aggregate interfaces.
 - The aggregate interfaces on the leaf nodes are shut down and then brought up.

201612170140

- Symptom: After the **ip forwarding-conversational-learning** command is executed, the switch does not issue a blackhole ARP entry when receiving unknown unicast traffic.
- Condition: This symptom might occur if the **ip forwarding-conversational-learning** command is executed.

201702270471

- Symptom: On an IGMP snooping-enabled VSI, residual multicast entries exist if more than 4 K multicast entries are generated.
- Condition: This symptom might occur if more than 4 K multicast entries are generated for an IGMP snooping-enabled VSI.

201703060189

- Symptom: In an EVPN-DCI network, after a VM migrates from ED 1 to ED 2, on ED 3 the route for reaching the VM points to ED 1.
- Condition: This symptom might occur if a VM migrates between two EDs in an EVPN-DCI network.

201702210531

- Symptom: A VXLAN tunnel interface cannot forward unicast traffic to another interface on the same card.
- Condition: This symptom might occur if the following conditions exist:
 - A VXLAN tunnel uses an aggregate interface as the outgoing interface.
 - The aggregation group has only one member port, and the member port is in Unselected state.

201701050265

- Symptom: Traffic forwarding failure occurs in multicast VLANs when certain operations are performed.
- Condition: This symptom might occur if the following operations are performed:
 - a. Send IPv6 traffic in multicast VLANs.

- b. Send membership reports and leave messages in sequence in the multicast VLANs.

201703010222

- Symptom: PFC becomes ineffective on a 10-GE interface when certain operations are performed.
- Condition: This symptom might occur if the following operations are performed:
 - a. Execute the **priority-flow-control** command in the view of a 10-GE interface that is installed with a GE transceiver module.
 - b. Save the running configuration and reboot the switch with the transceiver module installed.
 - c. Execute the **priority-flow-control** command in the view of the 10-GE interface and re-install the GE transceiver module.

201702210388

- Symptom: On an interface, an IPv6 ACL that filters packets requiring to be sent to the control plane for forwarding does not take effect.
- Condition: This symptom might occur if an IPv6 ACL is applied to an interface to filter packets that require to be sent to the control plane for forwarding.

201612270074

- Symptom: The **undo qos wred apply** command takes effect only after the running configuration is saved and the switch is rebooted.
- Condition: This symptom might occur if the **qos wred apply** and **undo qos wred apply** commands are executed in sequence in interface view.

201703070364

- Symptom: The **qinq ethernet-type service-tag** command does not take effect on an interface after the running configuration is saved and the switch is rebooted.
- Condition: This symptom might occur if the following operations are performed:
 - a. Execute the **qinq ethernet-type service-tag** command in interface view.
 - b. Save the running configuration and reboot the switch.

201702280503

- Symptom: When certain conditions exist, the console port stops responding for about ten minutes.
- Condition: This symptom might occur if the following conditions exist:
 - a. A large number of users perform MAC authentication, and the authentication server issues ACLs.
 - b. The **port link-mode route** command is execute on a Layer 2 interface enabled with MAC authentication.

201701160436

- Symptom: BGP neighbors are disconnected when the interface that provides the source IP address for VXLAN tunnels to the neighbors is shut down.
- Condition: This symptom might occur if the interface that provides the source IP address for VXLAN tunnels is shut down.

201703140311

- Symptom: The **display this** command does not display BFD configuration.
- Condition: This symptom might occur if the **bfd min-receive-interval 1000** and **display this** commands are executed in sequence in Layer 3 Ethernet subinterface view.

201701090159

- Symptom: After a large number of VPN instances are created and some VPN instances are deleted, the routes that the switch learns cannot reach the upper limit.
- Condition: This symptom might occur if the following operations are performed:
 - a. Create a large number of VPN instances.
 - b. Delete some of the VPN instances.

201703130545

- Symptom: A Layer 3 Ethernet subinterface cannot be pinged after certain operations are performed.
- Condition: This symptom might occur if the following operations are performed:
 - a. Create Layer 3 Ethernet interfaces A and B and Layer 3 Ethernet subinterfaces A.x and B.x.
 - b. Delete Layer 3 Ethernet subinterface B.x.
 - c. Ping Layer 3 Ethernet subinterface A.x.

201703030085

- Symptom: A basic ACL that is configured to match outgoing fragments matches all packets on an interface.
- Condition: This symptom might occur if a basic ACL is applied to an interface to match outgoing fragments.

201703060609

- Symptom: If the number of spanning tree instances reaches the upper limit, the **display erps** commands output information after a long delay.
- Condition: This symptom might occur if ERPS is configured, and the number of spanning tree instances reaches the upper limit.

201703100439

- Symptom: When a PBR policy is applied to a VLAN interface that provides the source IP address for a GRE tunnel, the PBR policy does not take effect, and the output interface to the next hop does not have traffic statistics.
- Condition: This symptom might occur if a PBR policy is applied to a VLAN interface that provides the source IP address for a GRE tunnel.

201703130281

- Symptom: If an unsupported action is added to a traffic behavior in a QoS policy applied to an aggregate interface, the add operation fails, and the traffic behavior becomes ineffective.
- Condition: This symptom might occur if the following operations are performed:
 - a. Apply a QoS policy to an aggregate interface.
 - b. Add an unsupported action to a traffic behavior in the QoS policy.

201612260345

- Symptom: When BGP is configured, BFD flapping occurs on an interface.
- Condition: This symptom might occur if the ACL type for the BGP process has been issued for another process.

201703020030

- Symptom: IRF links cannot come up if switches labeled with the following product codes use 40-GE ports or 100-GE ports for IRF connection:
 - 5940 32QSFP+ JH396A.
 - 5940 2-slot Switch JH397A.

- 5940 4-slot Switch JH398A.
- 5940 48XGT 6QSFP+ JH394A.
- 5940 48SFP+ 6QSFP+ JH395A.
- 5940 48XGT 6QSFP28 JH391A.
- 5940 48SFP+ 6QSFP28 JH390A.
- Condition: This symptom might occur if 5940 switches use 40-GE ports or 100-GE ports for IRF connection.

201704260375/201705310089

- Symptom: When certain operations are performed, the switch cannot communicate with a remote CE.
- Condition: This symptom might occur if the following operations are performed on the switch:
 - Create a Layer 2 aggregate interface.
 - Create an Ethernet service instance that uses the **encapsulation default** criterion on the Layer 2 aggregate interface.
 - Assign two or more interfaces to the aggregation group.

201703300325

- Symptom: After the controller issues a static route to a VPN instance through NETCONF, traffic forwarding fails within the VPN instance.
- Condition: This symptom might occur if the controller issues a static route to a VPN instance through NETCONF.

201703020655

- Symptom: When the MAC authentication offline detect timer expires, users in a voice VLAN cannot go offline if they use an OUI address as the source MAC address.
- Condition: This symptom might occur if the following operations are performed:
 - Enable MAC authentication globally and on a port.
 - Enable voice VLAN on the port.
 - Execute the **undo voice-vlan mode auto** command on the port.

201701040426

- Symptom: An interface might fail to establish a MACsec connection if the **mka enable** and **undo mka enable** commands are repeatedly executed on the interface.
- Condition: This symptom might occur if the **mka enable** and **undo mka enable** commands are repeatedly executed on an MACsec-capable interface.

201703020089

- Symptom: The MACsec connection on an interface might flap constantly if the **mka enable** command is executed on other MACsec-capable interfaces.
- Condition: This symptom might occur if the **mka enable** command is executed on MACsec-capable interfaces.

201703300105

- Symptom: A dynamic aggregate interface cannot forward Layer 2 traffic when Unselected member ports become Selected after some Selected member ports are removed.
- Condition: This symptom might occur if the following operations are performed:
 - Execute the **link-aggregation lacp traffic-redirect-notification enable** command in system view, and execute the **link-aggregation mode dynamic** command in aggregate interface view.

- b. Remove any number of Selected member ports from the dynamic aggregation group when it contains Unselected member ports.

201611100480

- Symptom: Interfaces cannot be configured as IRF physical interfaces after they are repeatedly associated with and disassociated from IRF ports.
- Condition: This symptom might occur if multiple interfaces are repeatedly associated with and disassociated from IRF ports.

201607270159

- Symptom: Service chain flow entries fail to be issued or do not take effect after being issued if the next hop of a VXLAN tunnel changes.
- Condition: This symptom might occur if the next hop of a VXLAN tunnel changes.

201704270501

- Symptom: The CLI does not respond if Neutron is enabled and interfaces are assigned to an aggregation group.
- Condition: This symptom might occur if the following operations are performed:
 - a. Enable Neutron.
 - b. Assign interfaces to an aggregation group.

201705100319

- Symptom: If certain operations are performed on an IRF fabric, the management Ethernet interface on the subordinate cannot be pinged after the management Ethernet interface on the master is shut down.
- Condition: This symptom might occur if the following operations are performed:
 - a. Assign IP addresses to the management Ethernet interfaces on the master and the subordinate.
 - b. Repeatedly shut down and bring up the management Ethernet interfaces.
 - c. Shut down the management Ethernet interface on the master.

201705170605

- Symptom: The switch reboots unexpectedly when a loopback interface is created and then shut down.
- Condition: This symptom might occur if a loopback interface is created and then shut down.

201705190229

- Symptom: Broadcast traffic forwarding fails when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - a. An AC is created on an interface of a PEX and is then removed.
 - b. Another interface of the PEX in the same VLAN as the interface in step a receives Layer 2 broadcast traffic.

201705240663

- Symptom: If process placement is repeatedly performed, VSI interfaces cannot be created.
- Condition: This symptom might occur if the following conditions exist:
 - a. EVPN is configured.
 - b. Process placement is repeatedly performed.
 - c. A large number of VSI interfaces are created.

201705310228

- Symptom: After the switch reboots, all physical interfaces become invisible at the CLI.
- Condition: This symptom might occur if the following operations are performed:
 - a. Use the **irf member renumber** command to modify the IRF member ID of the switch.
 - b. Upgrade the software to version F2607.

201701220392

- Symptom: On a two-chassis IRF fabric, the management Ethernet interfaces do not come up after they are installed with GE transceiver modules.
- Condition: This symptom might occur if GE transceiver modules are installed in management Ethernet interfaces of a two-chassis IRF fabric.

201705180551

- Symptom: The switch cannot communicate with directly connected peers if uRPF is enabled and then disabled.
- Condition: This symptom might occur if uRPF is enabled and then disabled.

201706060811

- Symptom: An Ethernet service instance on an aggregate interface is associated with a VSI. When new member ports join the corresponding aggregation group, the VSI cannot forward traffic.
- Condition: This symptom might occur if an Ethernet service instance on an aggregate interface is associated with a VSI, and new member ports join the corresponding aggregation group.

201704190367

- Symptom: The switch reboots unexpectedly if the **reset arp** and **display buffer usage** commands are executed in sequence multiple times.
- Condition: This symptom might occur if the **reset arp** and **display buffer usage** commands are executed in sequence multiple times.

201703130234

- Symptom: When a member port leaves an aggregation group, the configurations on the port are restored to the state before it joins the aggregation group.
- Condition: This symptom might occur if the following conditions exist:
 - a. A VM moves from aggregate interface A to aggregate interface B.
 - b. All member ports of aggregate interface A leave the aggregation group.

201703080626

- Symptom: When SoapUI is used to issue an xml configuration equivalent to **snmp-agent usm-user v3 user-role**, the maximum user role length is 62 bytes instead of 63 bytes.
- Condition: This symptom might occur if SoapUI is used to issue an xml configuration equivalent to **snmp-agent usm-user v3 user-role**.

201703310258

- Symptom: Member port A of a dynamic aggregation group cannot communicate with former member port B of the aggregation group.
- Condition: This symptom might occur if the following operations are performed:
 - a. Execute the **link-aggregation lacp traffic-redirect-notification enable** command in system view, and execute the **link-aggregation mode dynamic** command in aggregate interface view.
 - b. Assign port A to the dynamic aggregation group. Port A is Unselected.

- c. Remove Selected member port B from the dynamic aggregation group. Port A becomes Selected.

201703070096

- Symptom: The system prompts for resource insufficiency when the OSPF process is restarted.
- Condition: This symptom might occur if a node of a PBR policy contains two next hops for two ECMP routes.

201704190422

- Symptom: The **display qos policy control-plane** command does not display LACPDU statistics after a QoS policy is applied to the control plane to match LACPDU for traffic policing.
- Condition: This symptom might occur if a QoS policy is applied to the control plane to match LACPDU for traffic policing.

201705060379

- Symptom: When certain operations are performed, an EVPN VSI does not flood ARP packets.
- Condition: This symptom might occur if the following operations are performed on the switch:
 - a. Configure a VXLAN IP gateway group.
 - b. Execute the **flooding disable all** command on the EVPN VSI.
 - c. Save the configuration and reboot the switch.

201610310337

- Symptom: An SNMP tool fails to read the hh3cLswSubslotType and hh3cLswSubslotPortNum nodes.
- Condition: This symptom might occur if an SNMP tool is used to read the hh3cLswSubslotType and hh3cLswSubslotPortNum nodes.

201705170657

- Symptom: When both DRNI and FCoE are configured, the data link layer protocol of an FC interface cannot come up.
- Condition: This symptom might occur if both DRNI and FCoE are configured.

201705160022

- Symptom: A PBR policy is applied to two Layer 3 aggregation groups. When the PBR policy is removed for one aggregation group, it becomes ineffective on the other aggregation group.
- Condition: This symptom might occur if a PBR policy is applied to two Layer 3 aggregation groups, and the PBR policy is removed for one aggregation group.

201705190499

- Symptom: BFD sessions cannot come up after BFD is enabled for OSPF in GRE over IPv4 tunnel interface view.
- Condition: This symptom might occur if BFD is enabled for OSPF in GRE over IPv4 tunnel interface view.

201705150499

- Symptom: Unicast ARP packets are not sent to the ARP module of the switch.
- Condition: This symptom might occur if the following conditions exist:
 - EVPN is configured on the switch, and ARP flood suppression and proxy ARP are disabled.
 - An SDN controller issues an OpenFlow entry that matches ARP packets to the switch.

201705150541

- Symptom: An IRF fabric fails to forward traffic of a multicast VPN if the loopback group member interface for encapsulating private multicast traffic is not on the same slot as the public outgoing interface.
- Condition: This symptom might occur if the following conditions exist:
 - A multicast VPN is configured on an IRF fabric.
 - The loopback group member interface for encapsulating private multicast traffic is not on the same slot as the public outgoing interface.

201610120291

- Symptom: After the **placement reoptimize** command is executed, residual EVPN MAC address entries exist.
- Condition: This symptom might occur if the **placement reoptimize** command is executed on the switch configured with EVPN.

201704180198

- Symptom: After the **speed 100** command is executed on a 10-GE copper port on a switch labeled with one of the following product codes, the port cannot come up:
 - 5940 2-slot Switch JH397A.
 - 5940 4-slot Switch JH398A.
 - 5940 48XGT 6QSFP+ JH394A.
 - 5940 48SFP+ 6QSFP+ JH395A.
- Condition: This symptom might occur if the **speed 100** command is executed on a 10-GE copper port.

201704150167

- Symptom: The **apply precedence** command does not take effect if the **if-match vxlan-id** command has been configured for the same PBR policy.
- Condition: This symptom might occur if the **if-match vxlan-id** and **apply precedence** commands are configured for a PBR policy.

201703290346

- Symptom: A PBR policy cannot be applied to a Layer 3 aggregation group when certain operations are performed.
- Condition: This symptom might occur if the following operations are performed:
 - a. Set the mode of an Ethernet interface to Layer 3.
 - b. Create a Layer 3 aggregate interface.
 - c. Assign the Layer 3 Ethernet interface to the aggregation group.
 - d. Apply a PBR policy to the aggregation group.

201704190112

- Symptom: Incorrect ARP entries are learned on dynamic ACs, and traffic is forwarded incorrectly.
- Condition: This symptom might occur if dynamic ACs are configured on a DR member device.

201705250595

- Symptom: In an IRF 3.1 system, traffic forwarding between the parent fabric and PEXs fails after a master/subordinate switchover.
- Condition: This symptom might occur if the parent fabric contains both 5940 4-slot Switch JH398A and other models.

201705230578

- Symptom: In an IRF 3.1 system, broadcast traffic cannot be forwarded in a VSI if ACs are configured on two ports on a PEX and the two ACs are both mapped to the VSI.
- Condition: This symptom might occur if ACs are configured on two ports on a PEX and the two ACs are both mapped to the VSI.

201705240479

- Symptom: In an IRF 3 system, multicast traffic cannot be forwarded between a PEX and the parent device.
- Condition: This symptom might occur if the following conditions exist:
 - a. The 5930 switch with product code 5930-32QSPF+ JG726A, 5930-32QSFP+ TAA JG727A, 5930-2Slot+2QSFP+ JH178A, 5930-4Slot JH179A is used as a PEX.
 - b. One of the following switches is used as the parent device:
 - The 5940 switch with product code 5940 32QSFP+ JH396A ,5940 48XGT 6QSFP28 JH391A,5940 48SFP+ 6QSFP28 JH390A,5940 2-slot Switch JH397A or 5940 4-slot Switch JH398A

201704120628

- Symptom: An Ethernet service instance unexpectedly strips the VLAN tag of packets when forwarding the packets.
- Condition: This symptom might occur if the following conditions exist:
 - The Ethernet service instance and the PW through which the PE device connects to the peer PE device reside on different IRF member devices.
 - The Ethernet service instance is configured to match frames that are tagged with the specified outer 802.1Q VLAN ID.
 - The PVID of the Layer 2 Ethernet interface on which the Ethernet service instance is configured is the same as the specified outer 802.1Q VLAN ID.

201704210574

- Symptom: On an EVPN network, traffic cannot be forwarded through VXLAN tunnels because VXLAN tunnels are not correctly established.
- Condition: This symptom might occur if the following conditions exist:
 - Two VTEPs are virtualized by DRNI into a DR system and the DR system acts as a leaf node.
 - The leaf node first receives IP prefix routes advertised by a peer leaf node and then receives IMET routes from the same peer leaf node.

201704010366

- Symptom: An IRF subordinate device reboots repeatedly.
- Condition: This symptom might occur if the following operations are performed:
 - a. Set the IRF link down report delay to 0 milliseconds on the IRF fabric by using the **irf link-delay 0** command.
 - b. Save the configuration and then reboot the IRF fabric.

201704280079

- Symptom: An IRF fabric cannot be set up.
- Condition: This symptom might occur if a 40-GE interface inserted with a 100-GE cable is used as an IRF physical interface.

201705150181

- Symptom: An error of insufficient resources occurs when a rule is added to an existing ACL when only one ACL resource exists.
- Condition: This symptom might occur if a rule is added to an existing ACL when only one ACL resource exists.

201705230320

- Symptom: In an IRF 3.1 system, duplicated packets exist on a cascade port.
- Condition: This symptom might occur if unicast storm control is enabled by using the **storm-constrain unicast** command on a physical interface of the upstream port on the PEX.

201706030183

- Symptom: The values for the **EFP counter** field in the output from the **display qos-acl resource** command are incorrect.
- Condition: This symptom might occur if a large number of ACL rules used to count outbound traffic are configured.

201703060002

- Symptom: The **ExpectIP** attribute cannot be set to an all-F IPv6 address through NETCONF.
- Condition: This symptom might occur if the **ExpectIP** attribute is set to an all-F IPv6 address through NETCONF.

201703080615

- Symptom: The system mistakenly notifies that the **logon-page bind device-type** command in which the required option **file file-name** is not specified is successfully issued through NETCONF.
- Condition: This symptom might occur if the **logon-page bind device-type type-name** command is issued through NETCONF and the required option **file file-name** is not specified in this command.

201704100282

- Symptom: An Ethernet service instance is created on a Layer 2 aggregate interface, and a QoS policy is applied to an Ethernet service instance. After the Ethernet service instance is deleted, the QoS policy is still displayed in the output from the **display qos policy interface** command executed on the Layer 2 Ethernet interface.
- Condition: This symptom might occur if the following operations are performed:
 - a. Create an Ethernet service instance on a Layer 2 Ethernet interface.
 - b. Set the link type of the Layer 2 Ethernet interface to trunk.
 - c. Apply a QoS policy to the Ethernet service instance.
 - d. Delete the Ethernet service instance by using a controller.
 - e. Display the QoS policies applied to the Layer 2 Ethernet interface.

201703280584

- Symptom: If the controller is bound to multiple ACs, the switch fails to create ACLs issued by the controller.
- Condition: This symptom might occur if the controller is bound to multiple ACs.

201704140107

- Symptom: An interface has been assigned to an existing Layer 3 aggregation group, but the interface is not displayed as a member port of the Layer 3 aggregation group in the detailed information about the aggregation group.

- Condition: This symptom might occur if an interface is assigned to an existing Layer 3 aggregation group.

201704100093

- Symptom: The configuration file becomes very large after a large number of complicated ACs are configured.
- Condition: This symptom might occur if a large number of complicated ACs are configured.

201703220371

- Symptom: The network takes a long time to converge after an ISSU.
- Condition: This symptom might occur if both DRNI and ECMP are configured on the network and both incoming and outgoing Layer 3 traffic is forwarded on the network.

201703100235

- Symptom: Ports on the 5930 24p 10GBase-T and 2p QSFP+ Mod JH182A interface module of the 5940 2-slot Switch JH397A or 5940 4-slot Switch JH398A switch cannot forward traffic.
- Condition: This symptom might occur if the 5930 24p 10GBase-T and 2p QSFP+ Mod JH182A interface module is installed into the 5940 2-slot Switch JH397A or 5940 4-slot Switch JH398A switch and MACsec is enabled on the copper ports of the interface module.

201703200296

- Symptom: Layer 2 packets are duplicated when being forwarded out of a VXLAN tunnel interface.
- Condition: This symptom might occur if all ARP entries are cleared by using the **reset arp all** command after the VXLAN tunnel is established.

201704060276

- Symptom: After an interface module is removed from a PEX, ports on the removed interface module are still displayed as physical interfaces that can be used as member interfaces of the upstream port on the PEX in the **display system internal pex upstreamport** command output.
- Condition: This symptom might occur if an interface module is removed from a PEX.

201610290021

- Symptom: On an EVPN network, deleted EVPN MAC address entries still remain.
- Condition: This symptom might occur if the following operations are performed:
 - a. Send 32K Layer 2 unicast packets to a leaf node, and the source MAC address of each Layer 2 unicast packet is different.
 - b. Delete the VXLAN configuration on a VSI and then create another VXLAN on the VSI.
 - c. Repeat steps **a** and **b**.

201704280081

- Symptom: A QoS policy that uses an IPv6 ACL as a match criterion cannot be applied to interfaces after a rule in the IPv6 ACL is modified.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure a QoS policy that uses an IPv6 ACL as a match criterion.
 - b. Apply the QoS policy to interfaces of different types.
 - c. Modify a rule in the IPv6 ACL.

201705040449

- Symptom: The network-side many-to-one VLAN mappings fail to replace the SVLAN of downlink traffic with the CVLAN.

- Condition: This symptom might occur if many-to-one VLAN mappings are configured and downlink traffic is forwarded.

201705090236

- Symptom: On an IRF 3 system, Layer 3 traffic cannot be forwarded if the outgoing interface is a Layer 3 aggregate interface.
- Condition: This symptom might occur if the outgoing interface of Layer 3 traffic is a Layer 3 aggregate interface.

201704110633

- Symptom: A PBR policy still takes effect when the next hop for a destination IP address in the policy is inconsistent with the next hop for the destination IP address in the routing table.
- Condition: This symptom might occur if the next hop for a destination IP address in the policy becomes inconsistent with the next hop for the destination IP address in the routing table.

201705260484

- Symptom: In an IRF fabric, an FC aggregate interface cannot come up after a master/subordinate switchover.
- Condition: This symptom might occur if member ports in the corresponding FC aggregation group reside on subordinate devices and a master/subordinate switchover takes place.

201706050469

- Symptom: Forwarded traffic is doubled on the outgoing interface.
- Condition: This symptom might occur if a Layer 3 aggregation member interface is repeatedly shut down and brought up.

201703280348

- Symptom: In a VCF fabric, the master spine node cannot assign an IP address to a loopback interface on the switch (acting as a leaf node).
- Condition: This symptom might occur if the management Ethernet interface on the switch is associated with an MPLS L3VPN instance.

201610100229

- Symptom: Layer 3 Ethernet subinterfaces cannot be deleted.
- Condition: This symptom might occur if the following operations are performed:
 - a. Shut down multiple Layer 3 Ethernet subinterfaces by using the **shutdown** command.
 - b. Disable VLAN termination on the Layer 3 Ethernet subinterfaces by using the **undo vlan-type** command immediately after step a.
 - c. Delete the Layer 3 Ethernet subinterfaces by using the **undo interface** command.

201703270168

- Symptom: In an IRF 3.1 system, a PEX might fail to come online after the parent device reboots.
- Condition: This symptom might occur if the parent device reboots.

201703100674

- Symptom: On a VXLAN network, a memory leakage occurs on the device. The device fails to forward protocol packets or the device reboots.
- Condition: This symptom might occur if a large number of VXLAN tunnels are configured and the next hops of multicast routes are frequently changed.

201703300377

- Symptom: After a port isolation group is deleted, ports previously assigned to the group still cannot communicate with each other.
- Condition: This symptom might occur if an aggregate interface and its aggregation member ports are assigned to a port isolation group and then the port isolation group is deleted.

201705040697

- Symptom: A memory leakage occurs if the switch is configured to collect statistics for traffic on IPLs.
- Condition: This symptom might occur if the switch is configured to collect statistics traffic on IPLs.

201705120468

- Symptom: The VXLAN single-path detection feature is not available.
- Condition: This symptom might occur if the VXLAN single-path detection feature is used.

201705160745

- Symptom: The link-down event on a port causes another port running LACP or RRPP to go down and then come up.
- Condition: This symptom might occur if the 5940 2-slot Switch JH397A or 5940 4-slot Switch JH398A switch is installed with the 5930 24p SFP+ and 2p QSFP+ Mod w Msec JH181A, 5930 24p SFP+ and 2p QSFP+ Mod JH180A, or 5930 24p CP and 2p QSFP+ Mod JH184A interface module and a 10-GE SFP+ port on the interface module goes down.

201706050075

- Symptom: The flow mirroring configuration on a Layer 3 Ethernet interface and its subinterfaces cannot be removed when the GRE encapsulation format is used for mirrored packets.
- Condition: This symptom might occur if flow mirroring is configured on a Layer 3 Ethernet interface and its subinterfaces and the GRE encapsulation format is used for mirrored packets.

201706050329

- Symptom: A deleted default route still remains if the default route has multiple equal-cost routes and the hardware resource mode is set to 4 for the MAC address table, ARP/ND table, and routing table.
- Condition: This symptom might occur if the default route has multiple equal-cost routes and the hardware resource mode is set to 4 for the MAC address table, ARP/ND table, and routing table.

201706050447

- Symptom: The switch cannot forward unknown frames after the MAC learning limit on an interface is reached even though the switch is enabled to forward unknown frames received on the interface after the MAC learning limit on an interface is reached.
- Condition: This symptom might occur if the following conditions exist:
 - The MAC learning limit is set on an interface by using the **mac-address max-mac-count** command.
 - The switch is enabled to forward unknown frames received on the interface after the MAC learning limit is reached by using the **mac-address max-mac-count enable-forwarding** command.
 - The MAC learning limit on the interface is reached.

201706060692

- Symptom: In an IRF fabric, unknown multicast data still can be forwarded in a multicast VLAN when dropping unknown multicast data packets is enabled for the multicast VLAN.

- Condition: This symptom might occur if dropping unknown multicast data is enabled for the multicast VLAN in the IRF fabric.

201706070556

- Symptom: The global uRPF configuration gets lost after the IRF master device reboots.
- Condition: This symptom might occur if uRPF is enabled globally and then the IRF master device reboots.

201706060447

- Symptom: The LED for a QSFP28 port on the 5940 48XGT 6QSFP28 JH391A switch does not operate correctly.
- Condition: This symptom might occur if a transceiver module is inserted into the QSFP28 port.

201706030360

- Symptom: Unknown unicast packets are flooded in an IRF fabric.
- Condition: This symptom might occur if a large number of OpenFlow MAC address entries are configured on the IRF fabric and a master/subordinate switchover occurs.

201706020358

- Symptom: The value obtained from MIB node hh3cLswSubslotType is incorrect.
- Condition: This symptom might occur if MIB node hh3cLswSubslotType is read.

201706020245

- Symptom: On a DRNI network, the DR member devices cannot ping each other.
- Condition: This symptom might occur if any DR member device pings the IPP of the other DR member device.

201705260663

- Symptom: The LED for a port in up state does not operate correctly.
- Condition: This symptom might occur after the port is assigned to a service loopback group by using the **port service-loopback group group-id** command.

201705220406

- Symptom: A memory leakage occurs on the 5940 2-slot Switch JH397A or 5940 4-slot Switch JH398A switch after the 5930 24p SFP+ and 2p QSFP+ Mod w Msec JH181A or 5930 24p 10GBase-T and 2p QSFP+ Mod JH182A interface module is removed and then reinstalled.
- Condition: This symptom might occur if the following operations are performed:
 - a. Install the 5930 24p SFP+ and 2p QSFP+ Mod w Msec JH181A or 5930 24p 10GBase-T and 2p QSFP+ Mod JH182A interface module into the 5940 2-slot Switch JH397A or 5940 4-slot Switch JH398A switch.
 - b. Enable MKA on the 24 SFP+ ports of the 5930 24p SFP+ and 2p QSFP+ Mod w Msec JH181A interface module or enable MKA on the 24 10GBASE-T ports on the 5930 24p 10GBase-T and 2p QSFP+ Mod JH182A interface module.
 - c. Remove and then reinstall the 5930 24p SFP+ and 2p QSFP+ Mod w Msec JH181A or 5930 24p 10GBase-T and 2p QSFP+ Mod JH182A interface module.

201705160766

- Symptom: The issued ACLs do not take effect.
- Condition: This symptom might occur if ACLs are issued to the switch by a controller.

201705120500

- Symptom: On a VPLS network, VPLS packets get lost or are duplicated when they are forwarded.

- Condition: This symptom might occur if a cross-device aggregate interface is configured on the IRF fabric and a cross-card aggregate interface is configured as an AC.

201704120593

- Symptom: An error of insufficient resources occurs after the OSPF process restarts.
- Condition: This symptom might occur if the following operations are performed:
 - a. Create two PBR nodes and specify two next hops for each node by using the **apply next-hop** command.
 - b. Restart the OSPF process.

201704110307

- Symptom: A subordinate device fails to join the IRF fabric after the IRF fabric reboots.
- Condition: This symptom might occur with a low probability if the master device and the subordinate device run different software versions and the IRF fabric reboots.

201704100203

- Symptom: After the ARP fast update for MAC address moves is enabled, the system mistakenly generates a log notifying that the MAC address of the IPP moves.
- Condition: This symptom might occur after the IPP is specified on the switch and ARP fast update for MAC address moves is enabled.

201703180384

- Symptom: A memory leakage occurs in the MPLS TE process after the primary CRLSP fails and then recovers.
- Condition: This symptom might occur if CRLSP backup is configured and the primary CRLSP fails and then recovers.

201703100626

- Symptom: An ISSU fails to upgrade the software from a version earlier than F2606 to F2606 for an IRF fabric.
- Condition: This symptom might occur if an ISSU is performed to upgrade the software from a version earlier than F2606 to F2606 for an IRF fabric and then the **issu run switchover** command is executed to perform an ISSU switchover.

201702130356

- Symptom: An IRF subordinate device fails to restart up after a reboot.
- Condition: This symptom might occur after the following operations are performed:
 - a. Install multiple weak patches on the IRF subordinate device and commit software changes by using the **install commit** command.
 - b. Reboot the subordinate device.

201704010062

- Symptom: After a 40-GE interface is split into four 10-GE breakout interface, only one breakout interface is available.
- Condition: This symptom might occur if the following operations are performed:
 - a. Split a 40-GE interface installed with an adaptor module into four breakout interfaces and then combine the breakout interfaces into a 40-GE interface.
 - b. Replace the adaptor module with a transceiver module on the 40-GE interface.
 - c. Split the 40-GE interface into four breakout interfaces.

201607250189

- Symptom: A VM fails to communicate with a VXLAN IP gateway.

- Condition: This symptom might occur if the following operations are performed:
 - a. Configure an Ethernet service instance to match any frames that do not have an 802.1Q VLAN tag.
 - b. Map the Ethernet service instance to a VSI and set the access mode to VLAN for the Ethernet service instance.
 - c. Ping the VXLAN IP gateway from a VM in the VSI.

201704070029

- Symptom: Layer 3 traffic cannot be forwarded on a VXLAN network.
- Condition: This symptom might occur if Layer 3 Ethernet subinterfaces are used to establish a VXLAN tunnel between the VTEP and the VXLAN IP gateway.

201703010593

- Symptom: Users on a port can pass MAC authentication to come online when the access port is configured to deny packets from any VLAN.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure the port to deny packets from any VLAN.
 - b. Enable MAC authentication on the port.
 - c. Enable the voice VLAN feature on the port.
 - d. Configure the port to operate in manual voice VLAN assignment mode.

201703020670

- Symptom: The state of a MAC address entry in the output from the **display mac-address** command is incorrect. The deleted MAC address entry still remains.
- Condition: This symptom might occur if the following conditions exist:
 - A user comes online through a port after passing MAC authentication or 802.1X authentication.
 - If the port is configured with MAC authentication, after the voice VLAN feature is enabled and then disabled on the port, the user goes offline and then comes online and the port does not receive voice data.
 - If the port is configured with 802.1X authentication, after the voice VLAN feature is enabled on the port, the user is forcibly logged off. Then, the user comes online and the port does not receive voice data.

201705040345

- Symptom: The **shutdown** configuration of an IRF physical interface on the HPE 5940 2-port QSFP+ and 2-port QSFP28 Module JH409A interface module gets lost after the interface module reboots.
- Condition: This symptom might occur if the following operations are performed:
 - a. Use a 100-GE interface on the HPE 5940 2-port QSFP+ and 2-port QSFP28 Module JH409A interface module as an IRF physical interface.
 - b. Execute the **shutdown** command on the interface.
 - c. Reboot the interface module.

201704190465

- Symptom: The switch reboots unexpectedly.
- Condition: This symptom might occur if the speed of the connection interfaces on both the switch and the connected device is repeatedly changed between 1000 Mbps and 10000 Mbps at the same time.

201704170594

- Symptom: A 10-GE fiber port comes up and then goes down repeatedly if both the fiber port and the connected port operate at 1000 Mbps.
- Condition: This symptom might occur if a 10-GE cable is used to connect the fiber port and the other port and both the ports operate at 1000 Mbps.

201705220233

- Symptom: On an EVPN network, an IRF subordinate device reboots unexpectedly.
- Condition: This symptom might occur if two VTEPs form a DRNI system and VXLAN tunnels are established.

201705260273

- Symptom: An FC interface module takes 10 minutes to reboot after all Ethernet interfaces on the FC interface module are changed to FC interfaces.
- Condition: This symptom might occur after all Ethernet interfaces on the FC interface module are changed to FC interfaces.

201705230361

- Symptom: An FC aggregation member port cannot come up if the member port is repeatedly assigned to and then removed from the FC aggregation group.
- Condition: This symptom might occur if the member port is repeatedly assigned to and then removed from the FC aggregation group.

201705250358

- Symptom: The data link layer state of all member ports in an FC aggregation group cannot become UP if all member ports reside on subordinate devices of an IRF fabric.
- Condition: This symptom might occur if all member ports of the FC aggregation group reside on subordinate devices of an IRF fabric.

201705270497

- Symptom: The 5940 2-slot Switch JH397A or 5940 4-slot Switch JH398A switch reboots unexpectedly after an interface module that cannot be identified by the switch is installed and then removed.
- Condition: This symptom might occur an interface module that cannot be identified by the switch is installed and then removed.

201702140194

- Symptom: On a DRNI network, a member port of the IPP remains in MAD DOWN state and cannot come up.
- Condition: This symptom might occur if the following operations are performed:
 - a. Delete the IPP.
 - b. Initiate a primary/secondary switchover.
 - c. Roll back the software to an earlier version.

201703040088

- Symptom: The IRF 3.1 process reboots unexpectedly.
- Condition: This symptom might occur if the following operations are performed:
 - a. Create a Layer 2 extended-link aggregate group on a PEX.
 - b. Assign two interfaces to the Layer 2 extended-link aggregate group.
 - c. Remove one of the member ports from the Layer 2 extended-link aggregate interface.

201703090559

- Symptom: After the priority of a port in an MSTI is set, this configuration cannot be obtained from the STP/StlInterfaces table or the STP/STIs table.
- Condition: This symptom might occur if the priority of the port is set in an MSTI and the MSTI is not contained in the MST region to which the switch belongs.

201703140680

- Symptom: On a two-tier PEX stack, ports on the tier-1 PEX cannot be directly assigned to the cascade port on the tier-2 PEX after the cascade port on the tier-1 PEX is shut down and then brought up and the tier-1 PEX comes online again. Only after the cascade port on the tier-2 PEX is recreated, ports on the tier-1 PEX can be assigned to the recreated cascade port.
- Condition: This symptom might occur if the following operations are performed:
 - a. Disable PEX connection capability on the cascade port of the tier-2 PEX.
 - b. Shut down the cascade port on the tier-1 PEX and change the PEX group on the cascade port.
 - c. Bring up the cascade port on the tier-1 PEX so that the tier-1 PEX can come online.
 - d. Assign ports on the tier-1 PEX to the cascade port on the tier-1 PEX.

201703160414

- Symptom: If a VPNv4 route learned from the peer device and a VPNv4 route advertised by the public network have the same route prefix, the VPNv4 route advertised by the public still remains after being deleted.
- Condition: This symptom might occur if a VPNv4 route learned from the peer device and a VPNv4 route advertised by the public network have the same route prefix and the VPNv4 route advertised by the public network is deleted.

201703160426

- Symptom: On an EVPN DC network, Layer 3 traffic cannot be forwarded.
- Condition: This symptom might occur if the following operations are performed:
 - a. Delete all VSIs on the edge device and initiate an active/standby MPU switchover.
 - b. On the VTEP, manually soft-reset BGP sessions to the edge device.
 - c. Reconfigure the deleted VSIs on the edge device.

201705090430

- Symptom: On an MLAG network, the switch sends two copies of Packet-in messages to the controller.
- Condition: This symptom might occur if OpenFlow is configured on the MLAG network.

201703160476/201704100416

- Symptom: The switch reboots because of Watchdog timeout.
- Condition: This symptom occurs with a low probability if the switch runs for a long time.

201704060127

- Symptom: The Layer 3 Ethernet subinterface statistics obtained from MIB node ifHCInOctets/ifHCOctets are incorrect.
- Condition: This symptom might occur if the Layer 3 Ethernet subinterface statistics are obtained through SNMP from MIB node ifHCInOctets/ifHCOctets.

201705100215

- Symptom: The switch fails to learn the ARP entry for a connected VM. As a result, the VM cannot communicate with other VMs in the same VXLAN.

- Condition: This symptom might occur if ARP suppression is enabled on the VSI of the VXLAN by using the **arp suppression enable** command.

201704210158

- Symptom: The VCF controller fails to issue VSIs to the switch.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure VLAN-to-VXLAN mappings on the VCF controller when the VCF controller is not enabled to preconfigure VXLANs for the switch.
 - b. VMs come online.
 - c. Modify the VLAN-to-VXLAN mappings on the VCF fabric.

201704060507

- Symptom: On a VXLAN network, VMs fail to move.
- Condition: This symptom might occur if the following conditions exist:
 - The switch acts as a VTEP and is connected to a VMware controller.
 - The VMware controller issues MAC address entries to the switch.
 - MAC address entries are not updated after VMs move.

201702240509/201704100420

- Symptom: On a VXLAN network, the switch (acting as a VTEP) does not handle untagged VLAN packets based on the PVID of the incoming interface.
- Condition: This symptom might occur if the following operations are performed on the interface through which VMs access the switch:
 - a. Configure the Ethernet service instance to match frames tagged with the PVID of the interface.
 - b. Map the Ethernet service instance to the VSI to which the VMs belong.
 - c. Change the PVID of the interface.

201702210001

- Symptom: The BGP process fails to start up after the switch installed with the **5940-cmw710-system-weak-patch-r2509p02h02.bin** patch reboots.
- Condition: This symptom occurs with a low probability after the switch installed with the **5940-cmw710-system-weak-patch-r2509p02h02.bin** patch reboots.

201610190134

- Symptom: A VM fails to obtain an IP address through DHCP when flooding is disabled for a VSI by using the **flooding disable broadcast** command.
- Condition: This symptom might occur if flooding is disabled for a VSI by using the **flooding disable broadcast** command.

201702270416

- Symptom: On an EVPN network with distributed IP gateways, the switch that acts as the IP gateway mistakenly learns the source MAC address for a packet in a MAC address entry of the VSI of an L3 VXLAN after the incoming interface of the packet is changed.
- Condition: This symptom might occur if the incoming interface of the packet is changed.

201703170256

- Symptom: The switch generates error message "Rx/Tx failure detected between the CPU and switching chip." after an IRF fabric splits.
- Condition: This symptom might occur if the switch runs for a long time.

Resolved problems in F2606

201703220412

- Symptom: Software upgrade or device reboot causes traffic interruption for more than 1 second.
- Condition: This symptom might occur if the switch configured with NRNI and ECMP upgrades the software or reboots when Layer 3 traffic is being forwarded continuously.

201609300240

- Symptom: On the HPE 5940 2-port QSFP+ and 2-port QSFP28 Module JH409A interface module, an interface connected to another interface through a five-meter 40G QSFP cable cannot come up after the interface module that runs at 60 °C for a long time is power cycled.
- Condition: This symptom might occur with a low probability if the HPE 5940 2-port QSFP+ and 2-port QSFP28 Module JH409A interface module runs at 60 °C for a long time and is power cycled.

201704250206

- Symptom: After the action in the traffic behavior of a QoS policy is changed to flow mirroring, the flow mirroring action cannot be applied.
- Condition: This symptom might occur if the action is changed to flow mirroring when the QoS policy is already applied to the outgoing traffic of an interface.

201704110362

- Symptom: During automated deployment, the BFD MAD configuration is not correctly deployed. Specifically, the port supposed to be used for BFD MAD detection is used as an IRF physical interface.
- Condition: This symptom might occur if all configurations on the switch are cleared and the switch is automatically deployed.

201704010616

- Symptom: An IRF fabric splits when a dead loop occurs.
- Condition: This symptom might occur if multiple VPLS PWs exist on the IRF fabric and the public network interface is shut down and then brought up.

201704270571

- Symptom: Among four consecutive 10-GE interfaces, some are used as IRF physical interfaces and the others are installed with GE transceiver modules for connection. The interfaces installed with GE transceiver modules cannot come up.
- Condition: This symptom might occur if some interfaces in four consecutive 10-GE interfaces are used as IRF physical interfaces and the other interfaces are installed with GE transceiver modules for connection.

201704060293

- Symptom: On a DRNI network, the switch cannot timely update MAC address entries.
- Condition: This symptom might occur if the following tasks are performed:
 - a. Specify an interface on the switch as the IPP.
 - b. Change an interface in up state to operate in Layer 3 mode.
 - c. Delete MAC address entries on the DR peer by using the **undo mac-address** command.

201704190505

- Symptom: A 10-GE interface is displayed as a GE interface in a certain condition.

- Condition: This symptom might occur if the 10-GE interface is installed with a GE cable or transceiver module and the interface module where the interface resides is removed and then re-installed.

201704130272

- Symptom: The system prompts operation failure when the speed of a 10-GE interface is set to 100 Mbps.
- Condition: This symptom might occur if the speed of a 10-GE interface is set to 100 Mbps.

201704240047

- Symptom: A 40-GE member port of a dynamic aggregation group receives a forwarded known unicast packet again.
- Condition: This symptom might occur if the aggregation member port receives a known unicast packet after it is split into four breakout interfaces and then the four breakout interfaces are combined.

201704200102

- Symptom: An IRF fabric splits unexpectedly when the interface speeds of multiple interfaces are set and then restored to the default.
- Condition: This symptom might occur if the interface speeds of multiple interfaces are set and then restored to the default.

201704150015

- Symptom: It takes a long time for traffic to switch to the backup LSP when the primary LSP fails on an MPLS FRR network.
- Condition: This symptom might occur if the primary LSP fails.

201704140110

- Symptom: On an IRF fabric, a port cannot be assigned to a Layer 3 aggregation group.
- Condition: This symptom might occur if a port on the IRF fabric is assigned to a Layer 3 aggregation group.

201704270352

- Symptom: On a multicast VPN network, the switch replicates multiple copies of multicast packets if multiple ports are assigned to a service loopback group on the switch.
- Condition: This symptom might occur if multiple ports are assigned to a service loopback group on the switch.

Resolved problems in F2605

201607250457

- Symptom: When Ctrl+C is pressed during the execution of the **repeat** command on an interface, the status of the interface becomes incorrect because some configurations are missing.
- Condition: This symptom might occur if the following operations are performed:
 - a. Execute the **repeat** command in interface view.
 - b. Press Ctrl+C to stop the execution of the **repeat** command.

201701230034

- Symptom: A link aggregation group is configured on the Neutron-enabled switch, and no IP address is assigned to the management Ethernet interface. When the **shutdown** command is executed in aggregate interface view, the CLI stops responding.
- Condition: This symptom might occur if the following conditions exist:

- a. A link aggregation group is configured on the Neutron-enabled switch.
- b. No IP address is assigned to the management Ethernet interface.
- c. The **shutdown** command is executed to shut down the aggregate interface.

201611140329

- Symptom: When certain conditions exist, command execution might fail when the switch is accessed through the console port.
- Condition: This symptom might occur if the following conditions exist:
 - a. Puppet is configured on the switch.
 - b. The **third-part-process start name** command is executed multiple times for a process.
 - c. The **third-part-process stop** command is executed for the same process.

201610090321

- Symptom: When a PVST-enabled VLAN is deleted and the spanning tree process is restarted, status of the interfaces in another PVST-enabled VLAN becomes incorrect, and spanning tree protocol packets are flooded in the VLAN.
- Condition: This symptom might occur if the following conditions exist:
 - a. PVST is enabled on the switch.
 - b. A PVST-enabled VLAN is deleted.
 - c. The spanning tree process is restarted.

201702140057

- Symptom: When certain conditions exist, Layer 3 traffic is interrupted transiently on an IRF fabric.
- Condition: This symptom might occur if the following conditions exist:
 - a. As the ED of a VXLAN-DCI network, the IRF fabric establishes a VXLAN-DCI tunnel to a remote ED.
 - b. An ED receives 256 K of ARP packets, and the other ED receives traffic sourced from the unresolved IP addresses.
 - c. An IRF master/subordinate switchover occurs.

201701230495

- Symptom: The switch might reboot unexpectedly when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - The switch forwards IPv6 traffic through multiple ECMP routes.
 - sFlow collects IPv6 traffic statistics.

201703020022

- Symptom: On an IRF fabric, a PBR policy is applied to the outbound direction of a VXLAN tunnel interface. After the configuration is saved and the IRF fabric is rebooted, the subordinate member reboots unexpectedly.
- Condition: This symptom might occur if the following operations are performed:
 - a. Apply a PBR policy to the outbound direction of a VXLAN tunnel interface on an IRF fabric.
 - b. Save the configuration and reboot the IRF fabric.

201701140260

- Symptom: An S-channel created through automatic S-channel negotiation is deleted because the S-channel receives a large number of LLDP protocol packets. When LLDP traffic stops, the S-channel is re-established, but it cannot forward EVB data traffic.
- Condition: This symptom might occur if the following conditions exist:

- a. An S-channel created through automatic S-channel negotiation is deleted because the S-channel receives a large number of LLDP protocol packets.
- b. The S-channel is re-established after LLDP traffic stops.

201702080195

- Symptom: The OpenStack Neutron component fails to issue configurations to the switch when the switch has routes to the OpenStack platform and the management Ethernet interface is down.
- Condition: This symptom might occur if the switch has routes to the OpenStack platform and the management Ethernet interface is down.

201609130528

- Symptom: In a VCF fabric, the loopback interface of a leaf node and the loopback interface of a spine node might be assigned the same IP address.
- Condition: This symptom might occur if the following conditions exist:
 - a. After the VCF fabric is deployed, the Director automatically assigns IP addresses to loopback interfaces.
 - b. A leaf node and a spine node are rebooted.

201607210384

- Symptom: The L2VPN feature cannot be disabled when the standby MPU is booting.
- Condition: This symptom might occur if the following operations are performed:
 - a. Execute the **l2vpn enable** command in system view.
 - b. Save the configuration and reboot the switch.
 - c. Execute the **undo l2vpn enable** command when the standby MPU is booting.

201612140098

- Symptom: Two PBR policies with different node numbers have the same matching VXLAN ID and ACL. When the **undo apply next-hop** command is executed for the PBR policy with a smaller node number, the switch prompts for resource insufficiency.
- Condition: This symptom might occur if two PBR policies with different node numbers have the same matching VXLAN ID and ACL, and the **undo apply next-hop** command is executed for the PBR policy with a smaller node number.

201701030522

- Symptom: The switch fails to issue VXLAN AC configurations to the driver when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - a. Two PBR policies with different node numbers have the same configuration.
 - b. The PBR policies are applied to the outbound direction of a VXLAN tunnel interface.

201701120218

- Symptom: A manually shutdown Layer 3 Ethernet subinterface can perform Layer 3 forwarding.
- Condition: This symptom might occur if the **shutdown** command is executed in Layer 3 Ethernet subinterface view.

201611020220

- Symptom: Both 802.1X authentication and MAC authentication are enabled on an interface. When users access the network through MAC authentication, multicast traffic cannot be forwarded.
- Condition: This symptom might occur if both 802.1X authentication and MAC authentication are enabled on an interface, and users access the network through MAC authentication.

201612070313

- Symptom: When certain conditions exist, a default IPv6 route cannot be issued to a VPN instance. As a result, IPv6 traffic cannot be forwarded.
- Condition: This symptom might occur if the following conditions exist:
 - a. 2 K VPN instances have been created on the switch.
 - b. A new VPN instance is created and associated with a Layer 3 interface.
 - c. A default IPv6 route is issued to the VPN instance.

201609120496

- Symptom: When performing Layer 3 forwarding, a Layer 3 Ethernet subinterface modifies the 802.1p priority of packets based on the DSCP-802.1p priority map.
- Condition: This symptom might occur if the following conditions exist:
 - a. The switch is a border gateway in a VXLAN or EVPN network.
 - b. The **qos trust dscp** command is executed on a Layer 3 Ethernet interface.
 - c. A subinterface is created on the Layer 3 Ethernet interface.

201702130078

- Symptom: The system prompts for operation failure when the **undo jumboframe enable** and **jumboframe enable** commands are executed in sequence in interface view.
- Condition: This symptom might occur if the **undo jumboframe enable** and **jumboframe enable** commands are executed in sequence in interface view.

201702070508

- Symptom: When a VPN instance is associated with a local Layer 3 Ethernet subinterface and a remote Layer 3 Ethernet subinterface, the remote subinterface cannot be pinged from the local subinterface.
- Condition: This symptom might occur if a VPN instance is associated with a local Layer 3 Ethernet subinterface and a remote Layer 3 Ethernet subinterface.

201610080382

- Symptom: After the **cfp port-trigger rdi action shutdown** command is executed on an interface, the interface might go down when the remote peer is up.
- Condition: This symptom might occur if the following conditions exist:
 - a. The **cfp port-trigger rdi action shutdown** command is executed on an interface.
 - b. The remote peer of the interface goes down, and the **undo shutdown** command is used to bring up the local interface.
 - c. The remote peer comes up.

201701190195

- Symptom: The system displays that the static ARP entry limit is reached when the limit is not reached.
- Condition: This symptom might occur if the following conditions exist:
 - a. The static ARP entry limit is reached.
 - b. All static ARP entries are deleted.
 - c. Static ARP entries are added.

201702040031

- Symptom: When certain conditions exist, EVPN ARP flood suppression entries (displayed by using **display evpn route arp suppression**) are inconsistent with ARP flood suppression entries (displayed by using **display arp suppression vsi**).

- Condition: This symptom might occur if the following conditions exist:
 - a. The underlay network uses OSPF to advertise routes for VXLAN tunnel interfaces.
 - b. BFD is configured for OSPF.
 - c. BFD flapping causes frequent VXLAN tunnel status changes.

201701120323

- Symptom: A member device in an IRF fabric cannot operate correctly.
- Condition: This symptom occurs if the following conditions exist:
 - DHCP snooping is enabled on the IRF fabric.
 - DHCP snooping trusted ports are configured on the IRF fabric.
 - A large number of DHCP clients come online and go offline.

201608050201

- Symptom: An error message of "Error: Load failed! RCID(0xFFFFFFFF)" is displayed when the DSCP port priority on an interface is changed by using the **qos priority dscp priority-value** command.
- Condition: None.

201702160316

- Symptom: Aggregation member ports cannot become Selected if a member port is configured with the **speed** command and then with the **undo speed** command.
- Condition: This symptom might occur if a member port is configured with the **speed** command and then with the **undo speed** command.

201701170179

- Symptom: Two VFC interfaces operating in E mode are connected over a network. One VFC interface comes up rapidly after going down. However, the other VFC interface comes up slowly and responds slowly to FIP packets after going down
- Condition: None.

201611300307

- Symptom: The switch cannot forward traffic after multiple ECMP link switchovers.
- Condition: This symptom occurs if the switch reboots and causes ECMP link switchovers.

201702070510

- Symptom: A 10G copper port cannot come up after its link mode is changed from Layer 2 mode to Layer 3 mode by using the **port link-mode** command.
- Condition: None.

201612130470

- Symptom: After an interface is configured as a customer-side port, IPv4 routes and ARP entries fail to be issued.
- Condition: This symptom occurs if the following operations are performed:
 - Configure a VLAN interface as a customer-side port, and bind the VLAN interface to a VPN instance. Configure another VLAN interface in the same way. ARP packets are transmitted between the two VLAN interfaces.
 - Configure a VSI interface as a customer-side port, and bind the VSI interface to a VPN instance. Configure another VSI interface in the same way. ARP packets are transmitted between the two VSI interfaces.

201702070326

- Symptom: On tier-2 PEXs, Layer 2 traffic cannot be forwarded for all VLANs except for VLAN 1.

- Condition: None.

201703310540

- Symptom: PFC does not take effect on a Layer 3 aggregate subinterface.
- Condition: This symptom occurs if the Layer 3 aggregate subinterface is a tunnel interface and receives double-tagged packets.

201703310516

- Symptom: The reachability of a remote VM cannot be determined.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure an IRF fabric as a VTEP.
 - b. Configure an AC and a VXLAN tunnel interface on different member devices in the IRF fabric.
 - c. Test the connectivity between a local VM connected to the AC and a remote VM by using the **emulate-ping vxlan** command.

201703290614

- Symptom: A 40GE interface cannot come up after it is installed with a 40G-to-10G adapter.
- Condition: This symptom occurs with a low probability if the following operations are performed:
 - Add the 40GE interface installed with the adapter to an aggregation group.
 - Shut down and bring up the interface repeatedly.

201703280197

- Symptom: Packets from the underlay network to the overlay network cannot be forwarded on an EVPN network.
- Condition: This symptom occurs if the following conditions exist:
 - A Layer 3 Ethernet subinterface or Layer 3 aggregate subinterface is created and associated with a VPN instance.
 - The subinterface receives packets destined for VMs on the EVPN network.

201703200552

- Symptom: The de-encapsulated uplink traffic cannot be transmitted at the wire speed on an EVPN network.
- Condition: This symptom occurs if the switch automatically creates a VXLAN tunnel and the actual outgoing interface of the VXLAN tunnel is an aggregate interface.

201608100311

- Symptom: An AC on the switch fails to forwards packets on an EVPN network.
- Condition: This symptom occurs if the switch receives VXLAN packets on a Layer 3 Ethernet subinterface.

201703110222

- Symptom: Traffic cannot be forwarded on an EVPN network.
- Condition: This symptom occurs if the following conditions exist:
 - Routes for VXLAN tunnels have multiple equal-cost next hops and the routes for different VXLAN tunnels have the same equal-cost next hops.
 - One of the VXLAN tunnels is shut down.

201703020542/201703080176

- Symptom: Packet loss occurs on the switch.

- Condition: This symptom occurs if the switch operates in an FCoE mode (excluding the Transit mode) and receives a large number of FCoE and FC packets.

201701190556

- Symptom: MAC address entries for all VSIs cannot be fully displayed in the output from the **display l2vpn mac-address** command.
- Condition: This symptom occurs if more than 4096 VSIs are created, all the VSIs are in the UP state, and traffic flows in each VSI.

201702160623

- Symptom: Layer 3 traffic cannot be forwarded on an EVPN network in certain conditions.
- Condition: This symptom occurs if the following conditions exist:
 - The switch learns a large number of routing entries and the maximum number of ARP entries is reached.
 - Process placement policies are applied to optimize placement of BGP, L2VPN, and RIB processes by using the **placement reoptimize** command.

201701240198

- Symptom: OpenFlow entry issuing takes a long time.
- Condition: This symptom occurs if a large number of OpenFlow meter entries are issued.

201701240366

- Symptom: ACL resources occupied by OpenFlow entries are mistakenly counted as reserved resources in the output from the **display qos-acl resource** command.
- Condition: This symptom occurs if OpenFlow is configured on the switch and the switch receives OpenFlow entries from the controller.

201701220044

- Symptom: The switch is added to a VCF controller as a physical network element. However, OpenFlow instance configuration might not be completely deployed to the switch.
- Condition: None.

201609020275

- Symptom: Part of the AC configuration gets lost after the switch reboots.
- Condition: This symptom occurs if Neutron is enabled on the switch and the topology changes.

201701240343

- Symptom: A VSI cannot come up on a VXLAN-DCI network.
- Condition: This symptom occurs if the VSI is associated with an AC and a VXLAN-DCI tunnel.

201701130114

- Symptom: IPsec SA negotiations for some data flows fail and the data flows cannot be transmitted.
- Condition: This symptom occurs if multiple data flows trigger IKE SA negotiations simultaneously.

201611240020

- Symptom: The switch uses the MAC address learned for a VM to reply to an ARP request that is initiated from another local VM to request the MAC address of the local VM.
- Condition: This symptom occurs if ARP flood suppression is enabled by using the **arp suppression enable** command.

201611230520

- Symptom: On a DRNI network, VLAN mapping configuration for the aggregate interface specified as the IPP does not take effect.
- Condition: This symptom occurs if VLAN mapping is configured on the aggregate interface to remark VLAN tags of incoming packets.

201612100273

- Symptom: On a DRNI network, the aggregate interface specified as the IPP cannot come up correctly.
- Condition: This symptom occurs if DRNI MAD is configured and the aggregate interface is shut down and then brought up.

201703200560

- Symptom: Packet loss still occurs on an interface after PFC is enabled.
- Condition: This symptom occurs if the following conditions exist:
 - A 40-km-long single-mode optical fiber is inserted into a transceiver module with a transmission distance of 40 km on the interface.
 - Traffic congestion occurs.
 - PFC is enabled on the interface and the connection distance of+ the interface is set to 40 km.

201702160624

- Symptom: In an EVPN network with distributed gateways, Layer 3 traffic is interrupted for over 5 seconds after VM migration.
- Condition: This symptom might occur if the following conditions exist:
 - a. ARP flood suppression is enabled on EVPN distributed gateways.
 - b. VMs migrate between gateways.

201703110069

- Symptom: A memory leakage occurs during VM migrations.
- Condition: This symptom occurs if the following conditions exist:
 - The switch acts as a VTEP and the switch is enabled with MAC address move notifications and ARP fast update for MAC address moves.
 - VMs in a DC repeatedly migrate between DCs.

201612090171

- Symptom: The switch acts as an EVPN gateway in an EVPN-DCI network. When a VM migrates from the switch to another device, the MAC address entry and the ARP entry for the VM are not deleted completely on the switch.
- Condition: This symptom might occur if the following conditions exist:
 - a. ARP flood suppression is enabled on the switch.
 - b. A VM migrates from the switch to another device.

201702220062/201612280333

- Symptom: Multiple ECMP routes shared by VXLAN tunnels have the same next hop. When the next hop of some ECMP routes becomes unavailable, it takes a long period of time for traffic to be switched to another route.
- Condition: This symptom might occur if the following conditions exist:
 - a. Multiple ECMP routes shared by VXLAN tunnels have the same next hop.

- b. The next hop of some ECMP routes becomes unavailable because interfaces go down or devices are rebooted.

201701170211

- Symptom: A QoS policy fails to be applied.
- Condition: This symptom occurs if the OVSDB controller deploys a QoS policy that does not contain a DSCP marking action.

201702040031/201702040019

- Symptom: VMs might fail to come online because BGP, EVPN, and ARP flood suppression entries are inconsistent.
- Condition: This symptom occurs if the following operations are performed:
 - a. Enable ARP flood suppression in an EVPN network.
 - b. The BGP protocol frequently flaps. Tunnel interfaces repeatedly go down and come up.

201702040027

- Symptom: The ARP replies from a tunnel interface are dropped.
- Condition: This symptom occurs if the following operations are performed:
 - a. Enable ARP flood suppression on VTEPs in an EVPN network.
 - b. The VM attached to the device sends an ARP request to a VM attached to a remote leaf.
 - c. The ARP flood suppression entry of the VM is not changed locally.

201702040038

- Symptom: ARP/RARP broadcast packets fail to be sent.
- Condition: This symptom occurs if ARP flood suppression and proxy ARP are disabled in an EVPN network.

201701200298

- Symptom: When a MAC move occurs, the first ARP/RARP packet is not sent to the protocol stack.
- Condition: This symptom occurs if the following operations are performed:
 - a. Enable ARP flood suppression in an EVPN network.
 - b. The VM attached to the device moves from a VTEP to another VTEP. A MAC move occurs.

201702090449

- Symptom: When interfaces are batch shut down and then brought up or the device is rebooted, it takes about 40 seconds from the time when the first interface comes up to the time when the last interface comes up.
- Condition: This symptom occurs if all interfaces of a device with 10GE Base-T interfaces are correctly connected.

201702150299

- Symptom: The device might reboot unexpectedly.
- Condition: This symptom occurs with a low probability if the CPU sends a unicast IP packet and the destination IP address of the packet is deleted from the outgoing interface.

201702160635

- Symptom: An EVPN VTEP deletes the ARP flood suppression entry for a remote VM after the VM migrates between remote VTEPs for the first time.
- Condition: This symptom might occur if the following conditions exist:
 - a. A remote VM has been online for over 25 minutes.

- b. The remote VM migrates between remote VTEPs for the first time.

Resolved problems in F2604

201612270341

- Symptom: The system prompts that the memory is insufficient when SNMP is used to frequently read the BGP neighborhood information.
- Condition: This symptom occurs if SNMP is used to frequently read BGP information.

201612300048

- Symptom: An IRF fabric fails to forward Layer 3 traffic for VSIs in certain conditions.
- Condition: This symptom might occur if the following conditions exist:
 - a. The IRF fabric changes its bridge MAC address as soon as the address owner leaves, or the IRF bridge MAC persistence timer is set to 12 minutes.
 - b. TRILL is enabled on the IRF fabric.
 - c. An IRF master/subordinate switchover occurs.

201612240154

- Symptom: The switch cannot forward overlay traffic after a VXLAN VSI interface is disassociated from a VPN instance.
- Condition: This symptom might occur if the following operations are performed:
 - a. Create a VSI interface and associate it with a VPN instance.
 - b. Execute the **gateway vsi-interface** *vsi-interface-id* command to specify the VSI interface as the gateway interface of a VSI.
 - c. Disassociate the VSI interface from the VPN instance.

201612210081

- Symptom: A VXLAN tunnel cannot be disassociated from a VSI after the VXLAN tunnel is deleted.
- Condition: This symptom might occur if the following operations are performed:
 - a. Assign a VXLAN tunnel to a VSI.
 - b. Delete the VXLAN tunnel.
 - c. Disassociate the VXLAN tunnel from the VSI.

201612160537

- Symptom: In certain conditions, the switch cannot ping a directly connected peer.
- Condition: This symptom might occur if the following operations are performed:
 - a. Create a Layer 3 Ethernet subinterface and associate it with a VPN instance.
 - b. Connect the switch to the peer through the Layer 3 Ethernet subinterface.
 - c. Remove the interface card that hosts the main interface of the subinterface, and re-install the interface card.

201612160377

- Symptom: In l2gw mode, the switch reboots repeatedly after L2VPN is enabled.
- Condition: This symptom might occur if the following operations are performed:
 - a. Set the VXLAN hardware resource mode to l2gw.
 - b. Execute the **l2vpn enable** command to enable L2VPN.

201612150279

- Symptom: The controller fails to issue a security policy to the switch in certain conditions.
- Condition: This symptom might occur if the following conditions exist:
 - a. The switch and the VCF controller establish an aggregate link.
 - b. A security policy is created on the controller.
 - c. The controller issues the security policy to the switch.

201612090596/201612050219

- Symptom: In a VXLAN-DCI network, VM migration fails in certain conditions.
- Condition: This symptom might occur if the following conditions exist on the switch:
 - a. An Ethernet service instance is created on a Layer 2 aggregate interface and mapped to a VSI.
 - b. VMs access the VXLAN-DCI network through the Ethernet service instance.
 - c. The Ethernet service instance and related configuration are deleted.
 - d. The VMs migrate.

201612070066

- Symptom: When the switch acts as a VXLAN VTEP, ACL deployment fails in certain conditions.
- Condition: This symptom might occur if the following conditions exist:
 - a. Two nodes are configured for a PBR policy. The nodes have the same configuration and different node numbers.
 - b. The next hop of one node is removed.

201612060648

- Symptom: In certain conditions, OpenFlow entries cannot be issued to a subordinate member of an IRF fabric.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure OpenFlow on the IRF fabric.
 - b. Create an OpenFlow instance and deploy flow entries to the OpenFlow instance.
 - c. Save the configuration and reboot the IRF fabric or reboot a subordinate member.

201612060573

- Symptom: A VFC interface cannot forward FCoE traffic when it is bound to an aggregation member port.
- Condition: This symptom might occur if a VFC interface is bound to an aggregation member port.

201612020292

- Symptom: When the **tunnel all** command is executed for a VXLAN, manually created tunnels are not assigned to the VXLAN.
- Condition: This symptom might occur if the following operations are performed:
 - a. Create tunnels manually.
 - b. Execute the **tunnel all** command in VXLAN view.
 - c. Delete the tunnels and re-create them.

201611280438

- Symptom: EVPN is configured on the switch. After conversational remote MAC learning is enabled, blackhole MAC address entries are not issued when the switch receives unknown unicast traffic.

- Condition: This symptom might occur if the following conditions exist:
 - a. The **mac-address forwarding-conversational-learning** command is executed on the switch.
 - b. The switch receives unknown unicast traffic.

201611290605

- Symptom: An NMS cannot obtain information about the IldpXdot1dcbxRemPFCEnableEnabled MIB node.
- Condition: This symptom might occur if an NMS tries to obtain information about the IldpXdot1dcbxRemPFCEnableEnabled MIB node.

201611240304

- Symptom: In certain conditions, the VPN feature has traffic loss when a user logs in to the switch through 4G dialup.
- Condition: This symptom might occur if the switch is a P device enabled with GRE.

201611210633

- Symptom: When certain operations are performed, the interface management feature has memory leak.
- Condition: This symptom might occur if the following operations are repeatedly performed:
 - a. Create a Layer 3 Ethernet subinterface.
 - b. Set the mode of the subinterface to Layer 2.

20161110725

- Symptom: When the switch acts as a VTEP, a member port of a site-facing aggregate interface might receive traffic sent by itself.
- Condition: This symptom might occur if the following conditions exist:
 - a. An Ethernet service instance on a Layer 2 aggregate interface is mapped to a VSI.
 - b. The Layer 2 aggregate interface flaps.
 - c. The Ethernet service instance and related configuration are deleted.

201611180204

- Symptom: In certain conditions, the switch cannot forward VPLS traffic.
- Condition: This symptom might occur if the following conditions exist:
 - a. L3VPN and VPLS are configured on the switch.
 - b. An MPLS TE tunnel is configured as the public network tunnel for both L3VPN and VPLS.
 - c. Fast reroute is enabled for the MPLS TE tunnel.
 - d. Traffic of the MPLS TE tunnel is switched from its primary CRLSP to the bypass tunnel.

201611170345

- Symptom: When the switch acts as a VXLAN IP gateway, Tracert packets are dropped if local proxy ARP is disabled.
- Condition: This symptom might occur if local proxy ARP is disabled on the switch.

201611170240

- Symptom: When the number of Layer 3 Ethernet subinterfaces reaches the upper limit, some Layer 3 Ethernet subinterfaces cannot communicate with one another.
- Condition: This symptom might occur if the following conditions exist:
 - a. The number of Layer 3 Ethernet subinterfaces reaches the upper limit.

- b. The Layer 3 Ethernet subinterfaces created last communicate with the subinterfaces created earlier.

201611150224

- Symptom: The 5940 2-slot Switch JH397A5940 4-slot Switch JH398A switch might fail to display LLDP neighbor information for a copper port enabled with MACsec.
- Condition: This symptom might occur if the following conditions exist:
 - a. MACsec is enabled on a copper port of the 5940 2-slot Switch JH397A5940 4-slot Switch JH398A switch.
 - b. LLDP neighbor information is displayed for the copper port.

201611110455

- Symptom: In certain conditions, system ACL resources are occupied after the switch reboots without loading configuration.
- Condition: This symptom might occur if the following operations are performed:
 - a. Execute the **dhcp flood-protection enable** command on an interface.
 - b. Reboot the switch without any configuration.

201611140258/201611020637

- Symptom: In certain conditions, the VLAN module has slow memory leak.
- Condition: This symptom might occur if the following conditions exist:
 - a. A Layer 2 aggregate interface is configured as a trunk port and assigned to VLANs.
 - b. The Layer 2 aggregate interface is shut down and the brought up.

201611090374

- Symptom: When both 802.1X authentication and the EAD assistant feature are enabled, the EAD assistant feature does not take effect.
- Condition: This symptom might occur if the following conditions exist:
 - a. Both 802.1X authentication and the EAD assistant feature are enabled.
 - b. Users perform 802.1X authentication.

201611070471

- Symptom: The BGP process restarts unexpectedly when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - a. The number of BGP routes exceeds the upper limit.
 - b. BGP route flapping occurs constantly.

201610280470

- Symptom: A Layer 2 aggregate interface fails to forward broadcast traffic.
- Condition: This symptom might occur if a Layer 2 aggregate interface receives broadcast traffic.

201609280176

- Symptom: In certain conditions, PFC does not take effect on a 100-GE interface.
- Condition: This symptom might occur if the following conditions exist:
 - a. PFC is configured on the switch and the peer.
 - b. The output rate is limited to 6000 pps on the peer end.
 - c. The local end receives packets with a rate higher than 6000 pps.

201609130473/201608120534

- Symptom: On an IRF 3.1 fabric, configuration cannot be issued to interfaces on PEXs.

- Condition: This symptom might occur if the following conditions exist:
 - a. The parent fabric is rebooted.
 - b. The spanning tree feature is enabled and disabled globally for multiple times.
 - c. The **undo port trunk permit vlan all** command is executed on interfaces of PEXs.

201609120667

- Symptom: When L2VPN is disabled, VXLAN-DCI is not disabled correspondingly.
- Condition: This symptom might occur if the following conditions exist:
 - a. Execute the **dcu enable** command to enable VXLAN-DCI.
 - b. Disable L2VPN.

201612210431

- Symptom: In certain conditions, the switch displays deadloop log messages after a reboot.
- This symptom might occur if the following conditions exist:
 - a. The number of VLAN interfaces reaches the upper limit.
 - b. Kernel thread deadloop detection is enabled.
 - c. The configuration is saved and the switch is rebooted.

201612170046

- Symptom: In certain conditions, the switch cannot forward VXLAN traffic.
- Condition: This symptom might occur if the following operations are performed:
 - a. Execute the **l2vpn enable** command.
 - b. Create VSI-interface 1 and delete it.
 - c. Execute the **undo l2vpn enable** command.
 - d. Execute the **l2vpn enable** command, create VSI-interface 1, and configure VXLAN settings.

201612150574

- Symptom: In certain conditions, a private network tunnel is re-established when some L3VPN settings are modified.
- Condition: This symptom might occur if the following conditions exist:
 - a. MPLS L3VPN is configured on the switch and the peer. The switch and the peer establish a private network tunnel.
 - b. Execute the **undo vpn-target vpn-target&<1-8> export-extcommunity** command to delete the route targets of the corresponding VPN instance, and execute the **vpn-target vpn-target&<1-8> export-extcommunity** command to configure the same route targets for the VPN instance.

201611300548

- Symptom: In an EVPN network, VMs fail to move from a switch acting as a distributed EVPN gateway to another device.
- Condition: This symptom occurs if conversational remote MAC learning has been enabled by using the **mac-address forwarding-conversational-learning** command.

201610270530

- Symptom: The switch reboots because MAC authentication requests exhaust the memory resources.
- Condition: This symptom occurs if the following conditions exist:
 - MAC authentication is enabled globally and on interfaces.
 - The switch has many users online.

- The switch receives a large number of packets with unknown MAC addresses.

201610270390

- Symptom: Collection of aggregation group statistics is slow.
- Condition: This symptom occurs if the statistics are collected through MIB.

201609070088

- Symptom: In a VXLAN network, the radar detection results are inaccurate.
- Condition: This symptom occurs if local-first load sharing is enabled by using the **ip load-sharing local-first enable** command on an IRF fabric.

201609050389

- Symptom: The switch learns incorrect MAC addresses.
- Condition: This symptom occurs if LLDP is configured.

201608050200

- Symptom: An AC fails to forward packets.
- Condition: This symptom occurs if the following operations have been performed:
 - a. Configure VPLS.
 - b. Save the configuration and reboot the switch.

201607220401

- Symptom: The names of the following interface modules fail to be obtained through MIB:
 - 5930 8-port QSFP+ Module (JH183A).
 - 5930 8-port QSFP+ Module (JH181A).
 - 5930 24p 10GBase-T and 2p QSFP+ Mod (JH182A).
 - 5930 24p SFP+ and 2p QSFP+ Mod (JH180A).
- Condition: This symptom occurs if the names are obtained through MIB.

201607190364

- Symptom: In a VXLAN network, a site-facing interface can be assigned to an aggregation group.
- Condition: This symptom occurs if the following operations have been performed:
 - a. Specify the site-facing interface as a VTEP access port by using the **vtep access port** command.
 - b. Configure AC-related settings on the VTEP access port.
 - c. Delete the AC-related settings on the VTEP access port.

201702220233

- Symptom: An interface cannot come up.
- Condition: This symptom occurs if the interface connects to the peer interface by using a fiber-to-copper conversion module.

201702170474/201701240003

- Symptom: The CLI does not respond to input commands.
- Condition: This symptom occurs if the following operations are performed:
 - a. Enable MAC authentication globally and on an interface.
 - b. Use the server to assign authorization ACLs.
 - c. Change the operating mode of the interface when a large number of users come online.
 - d. Delete the MAC authentication configuration.

201702150447

- Symptom: Layer 2 or Layer 3 traffic forwarding between cards fails.
- Condition: This symptom occurs if the switch is connected to another device and local-first load sharing is disabled.

201702150239

- Symptom: FC/FCoE packets fail to be forwarded.
- Condition: This symptom occurs if the switch operates in Transit mode and is directly connected to an FC/FCoE-capable 3PAR storage device.

201702140100

- Symptom: The data link layer state of an FC interface cannot become UP.
- Condition: This symptom occurs if the FC interface operates in E mode.

201702140090

- Symptom: Layer 3 VXLAN traffic cannot be forwarded after TRILL is enabled globally.
- Condition: This symptom occurs if both TRILL and VXLAN are configured on the device.

201702130078

- Symptom: The error message "Operation failed" appears after an interface is configured with the **undo jumboframe enable** and **jumboframe enable** commands successively.
- Condition: None.

201608180168

- Symptom: The device does not discard a VXLAN packet that encapsulates an inner packet with TTL 0.
- Condition: None.

201702080474

- Symptom: The device assigns an incorrect MAC address to an interface. The MAC address is not in the reserved range of MAC addresses.
- Condition: None.

201701100505

- Symptom: The Python process fails to start up after an IRF master/subordinate switchover.
- Condition: This symptom occurs if the following conditions exist:
 - The IRF fabric connects to a VMware host.
 - The IRF fabric runs LLDP and the VMware host runs CDP.
 - The configuration is saved and the IRF fabric is restarted.

201612210125

- Symptom: Packet statistics does not take effect on an AC in a VXLAN network with IP gateways.
- Condition: This symptom occurs if the following conditions exist:
 - The device acts as a VXLAN IP gateway.
 - Packet statistics is enabled on the AC.
 - Packet statistics are displayed on the VSI interface or the AC after the VSI interface and the AC receive packets.

201612050503

- Symptom: LDP flaps and causes traffic interruption on an IRF fabric.

- Condition: This symptom occurs if the following operations are performed:
 - a. Configure LDP NSR and TCP on the IRF fabric.
 - b. Reboot the master IRF device.

201612200173

- Symptom: One file or directory on the device might be displayed as multiple ones.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure the **password-control enable** command on the device.
 - b. Multiple users perform file operations at the same time, for example, perform file/directory creation or deletion operations.
 - c. Execute the **dir** command to display the files or directories.

201701030339

- Symptom: The device displays a log message showing that "No enough hardware resource for MPLS." even when the hardware resources are enough.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure LDP to dynamically establish LSPs.
 - b. Configure OSPF to transmit route information.
 - c. Save the configuration, and reboot the IRF fabric or perform an IRF master/subordinate switchover.

201611040073

- Symptom: The BGP sessions between BGP peers on the IRF master member might go down.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure BGP NSR for the IRF fabric.
 - b. A subordinate member device fails and the IRF fabric splits. As a result, the subordinate member device becomes MAD Down.

201609050326/201609050325/201610210311

- Symptom: Using an IPv6 address to synchronize the time failed.
- Condition: This symptom might occur if NTP is enabled and an IPv6 address is used to synchronize the time.

201611150079

- Symptom: After an interface is installed with a GE transceiver module, the interface cannot come up.
- Condition: This symptom occurs if the following operations are performed:
 - a. Bind the interface to an IRF port, and then unbind the interface from the IRF port.
 - b. Install a GE transceiver module in the interface.

201612210218

- Symptom: STP loops might occur at a low probability.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure STP on an IRF fabric.
 - b. View the STP status after a master/subordinate switchover.

201612200493

- Symptom: The flood suppression state is displayed as failed in a VSI.
- Condition: This symptom occurs if the following operations are performed:

- a. Configure the centralized VXLAN IP gateway group on the device.
- b. Execute the **flooding disable unknown-multicast unknown-unicast** command in VSI view.
- c. Execute the **flooding disable unknown-multicast** command in VSI view.
- d. In probe view, execute the **display system internal overlay flooding vsi vsi-name** command to display the flood suppression state in a VSI.

201701090556

- Symptom: A short traffic interruption might occur.
- Condition: This symptom occurs with a low probability if the following conditions exist:
 - The switch acts as a VTEP in a VXLAN network.
 - The outgoing interface for a packet moves from a member device of an IRF fabric to another member device.

201611030342/201611250050

- Symptom: On an EVPN-enabled IRF fabric, an OpenFlow channel to the VCF controller changes to **Failed(Se)** state after an IRF master/subordinate switchover.
- Condition: This symptom might occur if the following conditions exist:
 - a. An IRF fabric sets up a connection to the VCF controller.
 - b. EVPN is configured and the **fail-open mode secure** command is executed on the IRF fabric.
 - c. An IRF master/subordinate switchover occurs.

Resolved problems in F2603

201612050510

- Symptom: Memory leaks occur on the device.
- Condition: This symptom occurs if the device receives a large number of ARP messages.

201611300053

- Symptom: Broadcast, multicast, and unknown unicast suppression bandwidth settings do not take effect for a short time and then take effect later.
- Condition: This symptom occurs if the following tasks are performed:
 - a. Set the broadcast, multicast, and unknown unicast suppression bandwidth by repeatedly using the **restrain { broadcast | multicast | unknown-unicast } bandwidth** command.
 - b. Restore one of the broadcast, multicast, and unknown unicast suppression bandwidth settings to the default by using the **undo restrain { broadcast | multicast | unknown-unicast }** command.

201611290137

- Symptom: Memory leaks occur on the Socket and LIPC modules.
- Condition: This symptom occurs if the following tasks are repeatedly performed:
 - a. Exclude the specified VLANs from the VLANs in which traffic is forwarded in the OpenFlow forwarding process by using the **openflow normal-forward vlan** command.
 - b. Cancel the above configuration by using the **undo openflow normal-forward vlan** command.

201611150347

- Symptom: On an EVPN network, BGP might select a wrong route after a VM quickly moves between three devices.

- Condition: This symptom occurs if the sequence number in the local MAC/IP route is the same as that in the MAC/IP route advertised by remote peers.

201611240295

- Symptom: The CPU usage of the device is high.
- Condition: This symptom occurs if the TCP MSS is set on an interface of the device by using the **tcp mss** command and the device receives a large number of fragmented packets.

201611240281

- Symptom: The Smart Link configuration is lost after the device recovers configuration.
- Condition: This symptom occurs if the following tasks are performed:
 - a. Configure Smart Link on an IRF fabric.
 - b. Restart an IRF member device.
 - c. Recover the configuration of the IRF member device by using an .mdb file when indexes of interfaces on the member device change.

201611240229

- Symptom: Portal users fail authentication and cannot come online.
- Condition: This symptom occurs if the following conditions exist:
 - a. The device uses a RADIUS server to perform authentication and accounting on portal users.
 - b. A large number of users go offline and other users try to come online simultaneously.

201611240225/201611220839

- Symptom: OSPF neighbor relationships cannot be established in a DCN network.
- Condition: This symptom occurs if the following conditions exist:
 - OSPF is enabled.
 - The .mdb file is deleted.
 - An IRF master/subordinate switchover occurs.

201611240213

- Symptom: After an IRF fabric is rebooted, a GRE tunnel is associated with a VPN instance different from the one associated before the reboot.
- Condition: This symptom might occur if the following operations are performed on an IRF fabric:
 - a. Create a GRE tunnel and associate it with a VPN instance by using the **ip binding vpn-instance** command.
 - b. Save the running configuration and reboot the IRF fabric.

201611220829

- Symptom: BFD session flapping occurs when the switch has more than 80 BFD sessions.
- Condition: This symptom might occur if the switch has more than 80 BFD sessions.

201611220826

- Symptom: On an IRF fabric, the master's console port is inaccessible if a user exits and then re-logs in.
- Condition: This symptom might occur if the ttymgr process is restarted.

201611190240/201611110207

- Symptom: A 5940 VXLAN IP gateway cannot forward QinQ traffic.
- Condition: This symptom might occur if the following conditions exist on the gateway:

- a. The **arp suppression enable** command is executed.
- b. QinQ traffic is received by an Ethernet service instance that uses the **encapsulation s-vid** *vlan-id* criterion.

201611170457

- Symptom: In an SDN EVPN network, VM 1 and VM 2 cannot communicate when they are connected to the same VTEP and are in the same subnet.
- Condition: This symptom might occur if the following conditions exist:
 - a. ARP flood suppression is enabled on the controller.
 - b. The ARP entry for VM 1 is deleted on the controller, and VM 2 sends an ARP request to obtain VM 1's MAC address.

201611140380

- Symptom: The reserved fields in the common header of PathErr messages sent by the switch are not reset.
- Condition: This symptom might occur if RSVP is enabled on the switch.

201611110604/TB201611110653/201610310018

- Symptom: ARP flood suppression entries are lost, or MAC address entry synchronization fails for VMs that migrate to the switch.
- Condition: This symptom might occur if ARP flood suppression is enabled on the switch, or VMs migrate to the switch.

201611100339

- Symptom: A 40-GE port cannot come up when installed with a 40GE-to-10GE adapter.
- Condition: This symptom might occur if a 40GE-to-10GE adapter is installed in a 40-GE port, and a 10-GE cable is connected to the adapter.

201611030479

- Symptom: The log buffer cannot record log messages after the system time is set back.
- Condition: This symptom might occur if the system time is set back.

201610260182

- Symptom: An IKE negotiation fails after multiple IKE key negotiations.
- Condition: This symptom might occur if IPsec is configured.

201611030393/201608220120

- Symptom: An SSH or Telnet user cannot log in when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - a. SYN Cookie is enabled.
 - b. The client is not directly connected to the switch.
 - c. The SSH or Telnet user uses an IPv6 address of the switch.

201611030385/201611010370/201610260481

- Symptom: The CLI does not respond after a user logs in through a management interface or console port when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - a. Password control is enabled.
 - b. A large number of users log in to the switch at the same time.

201611010418/201610260468

- Symptom: A Windows 10 user cannot input characters out of the ASCII code range 0x20 to 0x7F when configuring a description.
- Condition: This symptom might occur if a Windows 10 user inputs characters out of the ASCII code range 0x20 to 0x7F when configuring a description.

201611010170

- Symptom: MAC address learning fails for MAC address migration on aggregation member ports when the number of aggregate interfaces exceeds 824.
- Condition: This symptom might occur if the number of aggregate interfaces exceeds 824.

201610280009

- Symptom: An NMS cannot obtain the value of the entPhysicalSerialNum MIB node when the 5940 2-slot Switch JH397A/5940 4-slot Switch JH398A switch uses an extension interface card.
- Condition: This symptom might occur if the 5940 2-slot Switch JH397A/5940 4-slot Switch JH398A switch uses an extension interface card.

201610210315

- Symptom: A 10-GE port cannot come up when installed with a GE transceiver module.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure a 10-GE port as an IRF physical interface.
 - b. Remove the 10-GE port from its IRF port.
 - c. Install a GE transceiver module in the 10-GE port.

201610170395

- Symptom: The destination device sends an UPDATE message 10 seconds after a VM moves.
- Condition: This symptom occurs if the VM moves from an IRF fabric to another IRF fabric.

201610090389/201610130536

- Symptom: The device reboots unexpectedly.
- Condition: This symptom occurs if encrypted IPsec packets are fragmented on an intermediate link.

201610090373

- Symptom: A Layer 3 Ethernet subinterface cannot be deleted.
- Condition: This symptom occurs if the following tasks are performed:
 - a. Shut down the Layer 3 Ethernet subinterface by using the **shutdown** command.
 - b. Disable Dot1q termination on the Layer 3 Ethernet subinterface by using the **undo vlan-type dot1q vid *vlan-id-list*** command.
 - c. Delete the Layer 3 Ethernet subinterface.

201610090343/201610100252

- Symptom: ARP or ND entries might fail to be assigned and some traffic cannot be forwarded.
- Condition: This symptom occurs if a large number of ND entries are learned and a large number of ND entries age out.

201610090328

- Symptom: An SSH user cannot log in to the device even though the user enters the correct username and password.
- Condition: This symptom occurs if the following conditions exist:
 - A user logs in to the device through SSH.

- Password control is enabled on the device.
- The user enters wrong passwords for several times.

201609010473

- Symptom: The load sharing algorithm in the **display ip load-sharing path** command output and that in the **display ip load-sharing mode** command output are different.
- Condition: None.

201608300539

- Symptom: The later-applied ACL cannot filter outgoing packets.
- Condition: This symptom occurs if the following condition exist:
 - Apply an IPv4 ACL and an IPv6 ACL to filter outgoing packets on an interface.
 - The number of rules in the IPv4 ACL is greater than 256 and smaller than 512.
 - The IPv6 ACL contains the following rules:
 - **rule rule-id permit icmpv6**
 - **rule rule-id permit ipv6 source source-address**
 - **rule rule-id permit tcp destination destination-address destination-port eq xx**

201608260459

- Symptom: On an EVPN network with distributed gateways, VXLAN tunnels cannot be established because the BGP NSR feature cannot operate correctly. Traffic cannot be forwarded.
- Condition: This symptom occurs if the following condition exist:
 - A two-chassis IRF fabric acts as a leaf node.
 - A large number of VSIs and VSI interfaces are configured on the IRF fabric, and each VSI interface is associated with a VPN instance.
 - The IRF fabric triggers a memory usage alarm notification.
 - An IRF master/subordinate switchover occurs.

201607190097

- Symptom: The device might fail to record abnormal stacks after an NMI.
- Condition: This symptom occurs if a device reboot is interrupted by the hardware.

201607250099

- Symptom: The device does not respond to commands when it works with a server.
- Condition: This symptom occurs if the following conditions exist:
 - Durable queue creation is enabled on both the device and the server.
 - L2 agent and L3 agent are enabled and then disabled on the device.

201607130077/201607040494

- Symptom: The device reboots unexpectedly after an interface applied with a PBR policy is shut down.
- Condition: This symptom occurs if the following tasks are performed on the device:
 - Configure a VXLAN tunnel.
 - Apply a PBR policy on the outgoing VXLAN tunnel interface.
 - Shut down the outgoing VXLAN tunnel interface.

201610100268/201608300267

- Symptom: It takes a long time to install a patch on the master device of an IRF fabric.

- Condition: This symptom occurs if this patch is first installed on the master device rather than the subordinate devices.

201606290195/201607070414

- Symptom: The device reboots unexpectedly after checking the connectivity to an IPv4 address by sending large echo request messages.
- Condition: This symptom occurs if the following conditions exist:
 - ATK and ADVPN are configured on the public network.
 - ADVPN packets are fragmented.
 - Check the connectivity to an IPv4 address by using the **ping -s 8100** command.

201608100377

- Symptom: Host routes cannot be used for static route recursion.
- Condition: This symptom occurs if only host routes are specified for static route recursion in a static route (that is, specify the **recursive-lookup host-route** parameter when configuring a static route).

201607210452

- Symptom: When IPv4 IS-IS MTR and IPv6 IS-IS MTR are enabled, the switch cannot obtain routes from a Cisco NX9000 device.
- Condition: This symptom might occur if IPv4 IS-IS MTR and IPv6 IS-IS MTR are enabled, and the peer is a Cisco NX9000 device.

201607070398/201606280643

- Symptom: When a Telnet user uses an overlength username, the switch might reboot for memory exhaustion.
- Condition: This symptom might occur if a Telnet user uses an overlength username.

201604250059/201308080141

- Symptom: In an IRF fabric configured with OpenFlow, delay occurs when you display flow table information for an OpenFlow instance.
- Condition: This symptom occurs if a large number of VLANs are associated with the OpenFlow instance.
-

Resolved problems in R2509P02

201607180171

- Symptom: On an IRF fabric, a multidevice Layer 2 aggregate link cannot forward traffic or forwards duplicate traffic.
- Condition: This symptom might occur if the following conditions exist:
 - An IRF fabric connects to device A through a multidevice Layer 3 aggregate link and connects to device B through a multidevice Layer 2 aggregate link.
 - Device A sends multicast traffic to the IRF fabric.
 - The IRF fabric forwards the multicast traffic to device B through the Layer 2 aggregate link.

201609090145

- Symptom: In an IRF 3.1 system, the MAC address entries on devices cannot be displayed.
- Condition: This symptom occurs if the following operations are performed:
 - a. Use IMC to periodically obtain MAC addresses of the IRF 3.1 system.

- b. Execute the **display mac-address** command on the parent device.

201609080466

- Symptom: In an EVPN network, the IGMP snooping configuration on a device does not take effect.
- Condition: This symptom occurs if IGMP snooping is configured on VSI interfaces of the device in the EVPN network.

201609080154

- Symptom: Interface state flapping occurs on copper ports.
- Condition: This symptom occurs if the speed is set to 100 Mbps for the local copper port and 1Gbps/10Gbps on the peer copper port.

201609070399

- Symptom: When an aggregation group member interface in up state is configured as the management interface, the peer cannot be pinged.
- Condition: This symptom occurs if the management interface is not in Selected state.

201609070475/201609070465

- Symptom: Slow memory leak occurs on the device.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure the link aggregation management VLAN on the device.
 - b. The device receives traffic from the link aggregation management VLAN that needs to be forwarded through software.

201609070505/201609070454

- Symptom: Aggregate interfaces on an IRF fabric might fail to forward traffic.
- Condition: This symptom occurs if the Selected status of aggregation group member interfaces changes.

201609070023/201609070022

- Symptom: VSI configuration is lost.
- Condition: This symptom occurs if a VSI is created and packet statistics is enabled for the VSI when hardware resources are insufficient.

201609060223

- Symptom: A PBR policy still takes effect even if the match criteria configured for the PBR policy do not exist.
- Condition: This symptom occurs if the **if-match vxlan-id vxlan-id** command is configured in the PBR policy for the outbound packets of a VXLAN tunnel interface and the VXLAN specified in the command does not exist.

201608240137

- Symptom: A more layer of tags are encapsulated in GRE packets.
- Condition: This symptom occurs if the following operations are performed:
 - a. Establish a GRE tunnel between devices.
 - b. A service loopback group member interface receives inter-member device packets.

201609040010

- Symptom: The state of some OSPF neighbors returns to Init.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure an aggregation group that span IRF member devices and configure OSPF.

- b. Master/subordinate switchover occurs in the IRF fabric.

201609200007/201609200008

- Symptom: Slow memory leak occurs to the NETCONF OFP plug-in.
- Condition: This symptom occurs if a VCF controller is used to obtain the OpenFlow instance configuration of the switch.

201609190035/201609190037

- Symptom: After a 100-GE interface is removed from an IRF port and then bound to the IRF port, the interface goes down.
- Condition: This symptom occurs if a 100-GE interface is removed from an IRF port and then bound to the IRF port.

201609180328/201609180326

- Symptom: Some interfaces are lost.
- Condition: This symptom occurs if a 40-GE interface is split into four 10-GE breakout interfaces on a 5940 2-slot Switch.

201609080255/201609080257

- Symptom: A system-reserved ACL matches BFD_ARP_REPLY attack packets incorrectly.
- Condition: This symptom occurs if the following operations are performed:
 - a. BFD MAD is enabled on a VLAN interface on an IRF subordinate device.
 - b. The VLAN interface receives BFD_ARP_REPLY attack packets.

201609080100/201609080102

- Symptom: Reconfiguring port mirroring failed.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure more than four source ports for a mirroring group.
 - b. Delete the port mirroring configuration.
 - c. Reconfigure port mirroring.

201609070463/201609070467

- Symptom: The CLI on the IRF master device does not respond.
- Condition: This symptom occurs if the following operations are performed:
 - a. The device receives a large number of ARP packets.
 - b. Repeatedly execute the **display diagnostic-information** command and save the diagnosis information to files.

201609040029/201609040026

- Symptom: When a physical interface corresponding to a tunnel receives VXLAN packets, the physical interface cannot enqueue packets by the DSCP values of packets and the DSCP value of the packets is changed to 0.
- Condition: This symptom occurs if the physical interface corresponding to the tunnel is configured with the **qos trust dscp** command to trust the DSCP values of packets in a VXLAN or EVPN network.

201609020257

- Symptom: On a device operating in border mode, the bridge MAC address modified for an IRF fabric does not take effect.
- Condition: This symptom occurs if the **irf mac-address mac-address** command is used to modify the MAC address of the IRF fabric.

201608310372

- Symptom: In the inbound direction of an interface, applying ACLs to filter packets or applying a QoS policy that contain ACLs failed.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure IPv4 and IPv6 ACL rules.
 - b. The IPv4 ACL rule and the IPv6 ACL rule use the IPv4 quintuple and IPv6 quintuple to match packets, respectively.

201608260206

- Symptom: On an IRF fabric, an error occurs when the MAC address of a Layer 3 aggregate interfaces is deleted.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure the bridge MAC address for the IRF fabric, and configure a MAC address for the Layer 3 aggregate interface.
 - b. Delete the bridge MAC address configuration of the IRF fabric.

201608250214

- Symptom: When RESTful is used to deploy configuration, memory leak occurs to the xmlcfg process.
- Condition: This symptom occurs if HTTP-based RESTful is configured and the Post or Put operation is performed.

201608240391

- Symptom: After the master spine device is rebooted, exceptions occur to OSPF and BGP routes on leaf devices.
- Condition: This symptom occurs if the master spine device is upgraded and then rebooted in a network automatically deployed by VCF Fabric.

201608230350

- Symptom: After a 40-GE interface is split into four breakout interfaces, the breakout interfaces cannot be configured to operate in route mode.
- Condition: This symptom occurs if the following operations are repeatedly performed:
 - a. Split a 40-GE interface into four breakout interfaces.
 - b. Configure the breakout interfaces to operate in route mode.

201608220441

- Symptom: The CLI for an IRF subordinate device does not respond.
- Condition: This symptom occurs if an IRF fabric is split because the interface cards where the IRF physical interfaces reside are rebooted.

201608180323

- Symptom: Deleting traffic behaviors failed.
- Condition: This symptom occurs if the following operations are performed:
 - a. Configure 100 traffic classes and 100 traffic behaviors in a QoS policy.
 - b. Configure a flow mirroring action in a traffic behavior.
 - c. Apply the QoS policy to 10-GE breakout interfaces split from a 40-GE interface.
 - d. Combine the breakout interfaces, and delete the traffic behaviors in the QoS policy.

201608180217

- Symptom: IGMP snooping is enabled on the receiver-side device of a PIM-SM network. When interface flapping occurs on the device, Layer 2 multicast forwarding causes memory leaks.

- Condition: This symptom might occur if interface flapping occurs on the receiver-side device of a PIM-SM network.

201608060112

- Symptom: When strict uRPF check is enabled on the border gateway of an EVPN network, VMs cannot communicate with the external network.
- Condition: This symptom might occur if the following conditions exist:
 - a. The switch acts as the border gateway of an EVPN network.
 - b. The **ip urpf strict** command is executed to enable strict uRPF check on Layer 3 Ethernet subinterfaces of the switch.

201610240455

- Symptom: The switch is in a VXLAN network. When certain conditions exist, the switch reboots unexpectedly if it receives packets with two layers of VLAN tags.
- Condition: This symptom might occur if the **encapsulation s-vid vlan-id c-vid vlan-id-list** command is executed for an Ethernet service instance on the switch.

201610220110

- Symptom: ARP flood suppression does not take effect if the **hardware-resource vxlan l2gw** command has been executed.
- Condition: This symptom might occur if the **hardware-resource vxlan l2gw** command is executed.

201610200512

- Symptom: VXLAN traffic cannot be forwarded when Layer 3 aggregate subinterfaces are outgoing interfaces for traffic of VXLAN tunnels.
- Condition: This symptom might occur if Layer 3 aggregate subinterfaces are outgoing interfaces for traffic of VXLAN tunnels.

201610200210/201609290003

- Symptom: Members of an IRF fabric reboot unexpectedly when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - a. On the IRF fabric, ACs are configured on a large number of interfaces, and a large number of VLANs are configured.
 - b. The IRF bridge MAC address is configured.

201610180488

- Symptom: When multicast is enabled on the switch, multicast traffic loss occurs on Layer 3 subinterfaces.
- Condition: This symptom might occur if Layer 3 subinterfaces receive multicast traffic.

201610180217

- Symptom: Members of an IRF fabric reboot unexpectedly when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - a. The **placement reoptimize** command is executed on the IRF fabric.
 - b. The **mac-address forwarding-conversational-learning** command is executed.
 - c. The **undo mac-address forwarding-conversational-learning** command is executed.

201610170550

- Symptom: Interface speed modification causes memory leaks.
- Condition: This symptom might occur if the speed of an interface is modified.

201610170079

- Symptom: After ISSU incompatible upgrade is performed on subordinates of an IRF fabric, ISSU switchover cannot be performed on the master.
- Condition: This symptom might occur if ISSU incompatible upgrade is performed on subordinates of an IRF fabric.

201610130554

- Symptom: On an EVPN-DCI network, VSIs on 5940 EDs cannot come up.
- Condition: This symptom might occur if 5940 switches are used as EDs of an EVPN-DCI network.

201610130354

- Symptom: The **rule rule-id deny ip** rule of an ACL does not take effect.
- Condition: This symptom might occur if the following conditions exist:
 - a. The switch is a border gateway of a VXLAN network.
 - b. An ACL is applied to a VLAN interface to filter packets.
 - c. The **rule rule-id deny ip** rule is configured for the ACL by using the **rule [rule-id] { deny | permit } [{ { ipv4 | ipv6 | I2 | I4 } rule-string rule-mask offset } &<1-8>] [counting | time-range time-range-name] *** command.

201610130153

- Symptom: After certain operations, a queue scheduling profile cannot be applied to an interface.
- Condition: This symptom might occur if the following operations are performed:
 - a. Execute the **qos wrr queue-id group sp** command for a queue on the interface.
 - b. Execute the **undo qos wrr queue-id** command for the queue on the interface.
 - c. Execute the **qos qmprofile profile-name** command to create a queue scheduling profile that contains the configuration in step a.
 - d. Apply the queue scheduling profile to the interface.

201610130093

- Symptom: When an NMS reads the serial number of an interface card from the MIB, the serial number of the switch is returned.
- Condition: This symptom might occur if an NMS reads the serial number of an interface card from the MIB.

201610120080

- Symptom: Memory leaks occur when more than 500 VLAN interfaces are created on the switch.
- Condition: This symptom might occur if more than 500 VLAN interfaces are created on the switch.

201609280138/201601080493

- Symptom: An OpenFlow instance cannot be activated when it is configured to perform QinQ tagging for double-tagged packets passing an extensibility flow table.
- Condition: This symptom might occur if an OpenFlow instance is configured to perform QinQ tagging for double-tagged packets passing an extensibility flow table.

201609280042/201609280098

- Symptom: When more than 512 VLAN interfaces exist on the switch, Layer 3 Ethernet subinterfaces cannot be created.
- Condition: This symptom might occur if more than 512 VLAN interfaces exist on the switch.

201609270583

- Symptom: Memory leaks occur when aggregation groups are repeatedly created and deleted.
- Condition: This symptom might occur if aggregation groups are repeatedly created and deleted.

201609270183

- Symptom: When a .cfg configuration file that contains the L3 VXLAN ID configuration of the public instance is used for configuration restoration, the L3 VXLAN ID of the public instance is lost.
- Condition: This symptom might occur if the following operations are performed.
 - a. Create the public instance on the switch, and execute the **l3-vni vxlan-id** command in public instance view.
 - b. Save the configuration in a .cfg configuration file, and use the file for configuration restoration.

201609270125/201609270078

- Symptom: The switch reboots unexpectedly when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - a. The **ip forwarding-conversational-learning** command is executed.
 - b. Strict uRPF check is enabled when the switch is receiving packets.

201609230706

- Symptom: Errors occur when the switch floods multicast, broadcast, and unknown unicast traffic.
- Condition: This symptom might occur if the following conditions exist:
 - a. MLD snooping is enabled on a VSI.
 - b. An interface that hosts ACs for the VSI receives IPv6 multicast traffic.

201609230243

- Symptom: An IRF fabric cannot learn MAC addresses when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - a. MAC-based VLANs are configured on the IRF fabric, and the MAC learning limit is set on interfaces.
 - b. The MAC learning limit on interfaces is reached.

201609230082/201609230079

- Symptom: When tunnel interfaces of the switch are receiving bursts of ARP packets, some network segments of EVPN networks cannot be pinged.
- Condition: This symptom might occur if EVPN is configured on the switch, and tunnel interfaces of the switch are receiving bursts of ARP packets.

201609220666/201609220667

- Symptom: The globally-enabled uRPF feature becomes unavailable when aggregate interfaces receive packets.
- Condition: This symptom might occur if uRPF is globally enabled on the switch, and aggregate interfaces receive packets.

201609210174

- Symptom: When EVPN instances are frequently created and deleted, the switch reboots unexpectedly or stops responding.
- Condition: This symptom might occur if EVPN instances are frequently created and deleted.

201609190180

- Symptom: When certain conditions exist on an IRF fabric, BFD MAD is in faulty state, and the BFD process restarts unexpectedly.
- Condition: This symptom might occur if BFD MAD is configured on an IRF fabric and its peer, and the IRF fabric can receive BFD packets from the peer.

201608230424

- Symptom: Execution of the **ping -r** command fails.
- Condition: This symptom might occur if the **ping -r** command is executed.

201608170340

- Symptom: Execution of the **qos wfq { byte-count | weight }** command fails on a copper port after certain operations.
- Condition: This symptom might occur if the following operations are performed:
 - a. Execute the **qos sp** command on the copper port.
 - b. Execute the **qos wfq { byte-count | weight }** command on the copper port when the port is receiving traffic.

Resolved problems in R2509P01

None.

Resolved problems in R2509

201607230322

- Symptom: After certain operations, the switch cannot ping a directly connected peer through a Layer 3 aggregate interface.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure a Layer 3 aggregate interface.
 - b. Create 4094 subinterfaces on the Layer 3 aggregate interface and assign a MAC address to the Layer 3 aggregate interface.
 - c. Delete the MAC address for the Layer 3 aggregate interface and re-assign it a MAC address.

201607150047

- Symptom: An advanced ACL that contains a certain rule does not take effect.
- Condition: This symptom might occur if an advanced ACL contains one of the following rules:
 - **rule [rule-id] deny udp fragment source-port eq operator port1 destination-port eq operator port2**
 - **rule [rule-id] deny tcp fragment source-port eq operator port1 destination-port eq operator port2**

201607070522

- Symptom: AC configuration of a distributed VXLAN IP gateway is lost when certain conditions exist.
- Condition: This symptom might occur if a 5940 IRF fabric acts as the distributed VXLAN IP gateway, and an IRF master/subordinate switchover occurs.

201607070520

- Symptom: The **display link-aggregation verbose** command cannot correctly display the reference port information.
- Condition: This symptom might occur if the first bit of the switch's bridge MAC address is 0.

201607060264

- Symptom: RRPP configuration of an IRF fabric is lost when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - a. The IRF members start with .cfg files that contain RRPP configuration, and .mdb files are deleted after startup.
 - b. An IRF master/subordinate switchover occurs.

201607020149

- Symptom: A service chain flow entry issued by the controller does not take effect on a centralized VXLAN IP gateway when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - a. A 5940 IRF fabric acts as the centralized VXLAN IP gateway.
 - b. An OpenFlow instance is configured on the IRF fabric.
 - c. In OpenFlow instance view, the *local-port-number* field of the service chain flow entry specifies a port on the master.
 - d. An IRF master/subordinate switchover occurs.

201606280593

- Symptom: Layer 3 EVPN traffic cannot be forwarded after certain operations.
- Condition: This symptom might occur if the following operations are performed:
 - a. Delete VSIs and reconfigure VSIs.
 - b. Shut down and bring up physical interfaces that forward traffic of VXLAN tunnels.

201606280382

- Symptom: After certain operations, the CLI stops responding for a period of time, and the switch reboots unexpectedly.
- Condition: This symptom might occur if the following operations are performed:
 - a. Delete static ARP entries, and reconfigure the ARP entries for an interface by using the **arp static ip-address mac-address vlan-id interface-type interface-number** command.
 - b. In traffic behavior view, add an action of mirroring traffic to the interface by using the **mirror-to interface interface-type interface-number destination-ip destination-ip-address source-ip source-ip-address dscp dscp-value** command.
 - c. Repeatedly delete and create the traffic behavior.

201606270203

- Symptom: An interface cannot join an 802.1X guest VLAN when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - a. An 802.1X guest VLAN is configured by using the **dot1x guest-vlan guest-vlan-id** command on the interface.
 - b. Users on the interface pass 802.1X authentication.
 - c. The interface is shut down and then brought up.

201606150275

- Symptom: An interface goes down when its speed is set to 40 Gbps.

- Condition: This symptom might occur if the following operations are performed:
 - a. Connect a 100GE cable to an interface of a 5940 48XGT 6QSFP28/5940 48SFP+ 6QSFP28 switch.
 - b. Set the speed of the interface to 40 Gbps.

201606240436

- Symptom: A 5940 IRF fabric acts as a VXLAN IP gateway. After certain operations, a VSI interface on the gateway does not use the IRF bridge MAC address as expected.
- Condition: This symptom might occur if the following operations are performed:
 - a. Assign a MAC address to a VSI interface, and configure the IRF bridge MAC address.
 - b. Modify the VSI interface's MAC address and the IRF bridge MAC address.
 - c. Delete the MAC address of the VSI interface.

201606240418

- Symptom: When MAC-based VLAN is enabled and then disabled on an interface, online 802.1X users on the interface are forced off.
- Condition: This symptom might occur if MAC-based VLAN is enabled and then disabled on an interface.

201606170381

- Symptom: Two devices can establish a BFD session when they use different keys for BFD authentication.
- Condition: This symptom might occur if two devices use different keys for BFD authentication.

201606140140

- Symptom: A 5940 IRF fabric acts as a VXLAN VTEP. Members in the IRF fabric reboot repeatedly when they start up with certain configuration.
- Condition: This symptom might occur if 3000 VXLAN tunnels and 100 VSIs are configured on the IRF fabric.

201606130518

- Symptom: PEXs in an IRF 3.1 system fail to come online.
- Condition: This symptom occurs if the following conditions exist:
 - The IRF 3.1 system is rebooted repeatedly.
 - The cascade port on the parent device flaps.

201606070318

- Symptom: The CLI hangs.
- Condition: This symptom occurs if a PEX port group is deleted by using the **undo pex group** command on the parent device in an IRF 3.1 system.

201606070229

- Symptom: In an IRF 3.1 system, l2mc entries still exist after ports on a PEX are removed from all VLANs.
- Condition: This symptom occurs if the following tasks are performed:
 - a. Ports on the PEX are assigned to all VLANs by using the **port trunk permit vlan all** command.
 - b. Ports are removed from all VLANs by using the **undo port trunk permit vlan all** command.

201604260138

- Symptom: In an EVPN network, the device learns incorrect MAC address entries from the remote sites.
- Condition: This symptom occurs if the device receives MP-BGP routing information from remote sites.

201608250605

- Symptom: Layer 3 traffic cannot be forwarded in an EVPN network.
- Condition: This symptom occurs if the following conditions exist:
 - A Layer 2 aggregation group is created on devices in the EVPN network.
 - A non-reference member port in the aggregation group receives Layer 3 packets.

201608230326

- Symptom: Tags of ARP packets are removed when the packets are forwarded out of an IRF fabric.
- Condition: This symptom occurs if the following conditions exist:
 - a. An aggregate group is created on the IRF fabric.
 - b. An Ethernet service instance is created on the IRF fabric.
 - c. A VM connected to the IRF fabric receives ARP packets and sends ARP responses out of an aggregation member port.

201608230273

- Symptom: A VM cannot ping the centralized EVPN gateway in an EVPN network.
- Condition: This symptom occurs if the following tasks are performed:
 - a. Configure the device as an EVPN gateway.
 - b. Configure link aggregation on the device.
 - c. Configure an Ethernet service instance to match packets without VLAN tags.
 - d. Send ARP requests without VLAN tags from the VM.

201608230100

- Symptom: ACL configuration for known unicast packets is lost after an IRF fabric restarts.
- Condition: This symptom occurs if the following conditions exist on the IRF fabric:
 - Cross-device link aggregation is configured.
 - Link aggregation management VLANs are specified.
 - A port is configured as a management port for the aggregation group.

201608220102

- Symptom: The aggregation protocol flaps continuously.
- Condition: This symptom occurs if the following conditions exist:
 - Dynamic link aggregation is configured.
 - Link aggregation management VLANs are specified.
 - An aggregation member port is configured as a management port for its aggregation group.

201608190050

- Symptom: Traffic cannot be forwarded between Ethernet service instances in an EVPN network.
- Condition: This symptom occurs if the following tasks are performed on the 5940 device:
 - a. Create a VSI and an Ethernet service instance.

- b. Map the Ethernet service instance to the VSI.
- c. Delete configuration of the Ethernet service instance and that of the VSI.
- d. Create the same VSI and the same Ethernet service instance.
- e. Map the Ethernet service instance to the VSI.

201608180224

- Symptom: OpenFlow fails to deploy AC configuration.
- Condition: This symptom occurs if the following tasks are performed:
 - a. Establish a connection between the device and the OpenFlow controller.
 - b. Install a patch on the device and then reboot the device.

201608050481

- Symptom: The status of an aggregation member port on the device is **block**.
- Condition: This symptom occurs if DLDP is enabled on the aggregation member port by using the **dldp enable** command.

201608040353

- Symptom: The device reboots unexpectedly.
- Condition: This symptom occurs if a GRE tunnel interface is created on the device and VPN is configured in the view of the GRE tunnel interface.

201608020443

- Symptom: The device reboots unexpectedly.
- Condition: This symptom occurs if the following conditions exist:
 - a. Link aggregation management VLANs are specified.
 - b. An aggregate interface is created.
 - c. An aggregation member port is specified as the management port of the aggregation group.
 - d. The aggregate interface receives a large number of ARP requests.

201608020421

- Symptom: The CPU usage of the parent device in an IRF 3.1 system displayed by the **display cpu** command is high.
- Condition: This symptom occurs if the parent device has multiple PEXs connected.

201607280035

- Symptom: A VM cannot ping an EVPN gateway in an EVPN network.
- Condition: This symptom occurs if EVPN configuration is deployed through Neutron.

201607120173

- Symptom: Device information for PEXs in an IRF 3.1 system cannot be displayed by using the **display device** command.
- Condition: This symptom occurs if the cascade port of the parent device is shut down by using the **shutdown** command.

201607080440

- Symptom: The device reboot unexpectedly after the peer port is shut down by using the **shutdown** command.
- Condition: This symptom occurs if the following tasks are performed on a congested port:
 - a. Limit the rate of outgoing packets to a relatively low value.
 - b. Enable or disable PFC for an 802.1p priority value.

201607070563

- Symptom: In a VXLAN network with a centralized IP gateway, OpenFlow deploys flow entries containing different meter entries to the device. However, these flow entries fail to implement varied rate limiting.
- Condition: None.

201607070408

- Symptom: PBR does not take effect after the device performs a software upgrade.
- Condition: This symptom occurs if the device software is upgraded from Comware V5 to Comware V7.

201607060328

- Symptom: The LLDP process exits exceptionally after IRF physical interfaces are shut down and then brought up in bulk.
- Condition: This symptom occurs if the following tasks are performed:
 - a. Shut down IRF physical interfaces in bulk by using the **shutdown** command in interface range view.
 - b. Assign the IRF physical interfaces to an IRF port.
 - c. Bring up the IRF physical interfaces in bulk by using the **undo shutdown** command in interface range view.

201607050560

- Symptom: A Layer 3 aggregate subinterface on an IRF fabric cannot receive or forward packets after a master/subordinate switchover occurs on the IRF fabric.
- Condition: This symptom occurs if the following tasks are performed:
 - a. Configure a Layer 3 aggregate subinterface on the subordinate device in the IRF fabric.
 - b. Save the configuration and initiate a master/subordinate switchover.

201606280491

- Symptom: BFD flaps on an IRF fabric that has 32 BFD sessions established.
- Condition: This symptom occurs if SP queuing is enabled and then disabled on more than eight interfaces that do not have BFD sessions established.

201606280205

- Symptom: An IRF fabric splits after an ACL match criterion for an IPv6 node of an IPv6 PBR policy is changed.
- Condition: None.

201606230238

- Symptom: STP flaps in a PVST network.
- Condition: This symptom occurs if the following conditions exist:
 - VLANs are created in bulk on the IRF fabric.
 - The device receives a large number of Layer 2 broadcast packets with different source MAC addresses.

201605040410

- Symptom: Layer 3 traffic cannot be forwarded on an IRF fabric.
- Condition: This symptom occurs if a 40 GE interface on the master device in the IRF fabric is split into four 10 GE breakout interfaces.

201603240004

- Symptom: The device reboots unexpectedly.
- Condition: This symptom occurs if the device acts as the centralized EVPN gateway and L3 agent is enabled or disabled after the version of Neutron is updated.

201604260454

- Symptom: In a VXLAN network, the DHCP server receives duplicated DHCP requests.
- Condition: This symptom occurs if the DHCP relay agent is enabled on both a distributed VXLAN IP gateway and the border gateway.

201606060209

- Symptom: In an IRF fabric, traffic cannot be correctly forwarded after a patch is installed.
- Condition: This symptom occurs if the following conditions exist:
 - a. The device has a hot patch installed to fix STP problems.
 - b. The spanning tree protocol operates in PVST mode on the device.
 - c. VLANs have been irregularly added and deleted on the device.

201606300317/201606270528

- Symptom: When a Telnet user uses an overlength username, the switch might reboot for memory exhaustion.
- Condition: This symptom might occur if a Telnet user uses an overlength username.

201606230194/201606230190

- Symptom: On an IRF fabric, the **display mac-address** command does not display the MAC addresses learned on an aggregate interface.
- Condition: This symptom might occur if the following conditions exist:
 - A multichassis aggregate interface is configured.
 - Traffic of the aggregate interface is forwarded by only one IRF member.

201607160186

- Symptom: A VTEP cannot establish automatic VXLAN tunnels with a centralized VXLAN IP gateway group.
- Condition: This symptom might occur if the **vtep group remote** command specifies eight members for a centralized VXLAN IP gateway group on a VTEP.
-

Resolved problems in R2508

201607060206

- Symptom: A static MAC address entry cannot be configured on a Layer 2 aggregate interface.
- Condition: This symptom occurs if the following tasks are performed:
 - a. Create a Layer 2 aggregation group and configure a static MAC address entry on the Layer 2 aggregate interface.
 - b. Delete the Layer 2 aggregation group.
 - c. Create the same Layer 2 aggregation group and configure the same static MAC address entry on the Layer 2 aggregate interface.

201607130166

- Symptom: A host cannot ping a VSI interface or a loopback interface associated with the VPN instance of an L3VNI on the centralized EVPN gateway.
- Condition: This symptom occurs if the device acts as the centralized EVPN gateway.

201607130009

- Symptom: A VM connected to a VTEP cannot ping a loopback interface on the master device in a two-chassis IRF fabric.
- Condition: This symptom occurs if the following conditions exist:
 - A two-chassis IRF fabric acts as the VXLAN IP gateway.
 - A VTEP is connected to the subordinate device in the IRF fabric.
 - The outgoing port for VXLAN packets exchanged between the VTEP and the VXLAN IP gateway is an access port.

201607010268

- Symptom: The device cannot ping a Layer 3 aggregate subinterface on the peer device.
- Condition: This symptom occurs if the following tasks are performed:
 - a. Configure a Layer 3 aggregation group on the device and the peer device.
 - b. Create a Layer 3 aggregate subinterface on the peer device.
 - c. Split a 40-GE aggregation member port in the Layer 3 aggregation group on the device.

201606300512

- Symptom: Flow entries do not take effect on an interface that hosts an AC when certain conditions exist.
- Condition: This symptom might occur if the following conditions exist:
 - a. The AC uses the **encapsulation untagged** criterion or the **encapsulation s-vid** criterion that matches the PVID of the interface that hosts the AC.
 - b. VMs send untagged packets to the switch.

201606300296

- Symptom: A VXLAN tunnel cannot be established on an EVPN network.
- Condition: This symptom occurs if 5940 switches and Cisco devices exist in the EVPN network.

201606290422/201606200288

- Symptom: A Layer 3 interface broadcasts an ARP request whose source MAC address is all 0s to all the other Layer 3 interfaces.
- Condition: This symptom occurs if a Layer 3 interface receives an ARP request whose source MAC address is all 0s.

201606280096

- Symptom: Some traffic cannot be correctly forwarded by the border device on an EVPN network.
- Condition: This symptom occurs if the following tasks are performed:
 - a. Configure an aggregation group on the border device and a VTEP.
 - b. Shut down the aggregate interface by using the **shutdown** command.
 - c. Bring up the aggregate interface by using the **undo shutdown** command.

201606270576

- Symptom: A 100-GE interface cannot be shut down.
- Condition: This symptom occurs if the following tasks are performed:

- a. Bind a 100-GE interface on the 5940 48XGT 6QSFP28/5940 48SFP+ 6QSFP28 device to an IRF port and then remove the 100-GE interface from the IRF port.
- b. Shut down and then bring up the 100-GE interface, and repeat the operation.

201606270342

- Symptom: ACL resources in the inbound direction cannot be released.
- Condition: This symptom occurs if the following tasks are performed on a VLAN interface:
 - a. Enable IPv6SG and verify the source IPv6 address and MAC address for dynamic IPv6SG by using the **ipv6 verify source ip-address mac-address** command.
 - b. Disable IPv6SG by using the **undo ipv6 verify source** command.

201606240272

- Symptom: The CLI is stuck.
- Condition: This symptom occurs if a 40-GE interface is split into 10-GE breakout interfaces and then the 10-GE breakout interfaces are combined.

201606230053

- Symptom: Traffic cannot be forwarded on a distributed EVPN gateway.
- Condition: This symptom occurs if the following conditions exist:
 - A two-chassis IRF fabric acts as a distributed EVPN gateway.
 - A master/subordinate switchover occurs in the IRF fabric.

201606210498

- Symptom: The device reboots unexpectedly.
- Condition: This symptom occurs if the following conditions exist:
 - The device operates in PEX mode.
 - A physical interface on the device is bound to an IRF port.

201606170275

- Symptom: When rules are dynamically added to an ACL, the system prompts that the ACL resources are insufficient, and all rules in the ACL fails to be applied.
- Condition: This symptom occurs if the following conditions exist:
 - The ACL are applied to multiple ports.
 - Rules are dynamically added to the ACL.

201606170060

- Symptom: Packet filtering configuration on a Layer 3 aggregate interface or that on a member port in the aggregation group, whichever is earlier, takes effect.
- Condition: This symptom occurs if packet filtering is configured both on a Layer 3 aggregate interface and a member port in the aggregation group.

201606160166

- Symptom: After an ACL is removed from an interface, the ACL resources cannot be released.
- Condition: This symptom occurs if the following tasks are performed:
 - a. Configure an interface to filter outgoing packets.
 - b. Create an interface range by using the **interface range** command.
 - c. Add the interface in step a to the interface range.
 - d. Apply a user-defined ACL to filter packets on the interface range.
 - e. Remove the user-defined ACL from the interface range.

201606150402

- Symptom: The device reboots unexpectedly.
- Condition: This symptom occurs if the following tasks are performed:
 - a. Configure the device to operate in advance mode.
 - b. Apply an IPv4 ACL to filter packets and a QoS policy on an interface.
 - c. Apply an IPv6 ACL to filter packets on the interface.
 - d. Remove the IPv4 ACL, the IPv6 ACL, and the QoS policy on the interface.
 - e. Apply an IPv6 ACL to filter packets on the interface.
 - f. Apply an IPv4 ACL to filter packets on the interface.
 - g. Remove the IPv4 ACL and the IPv6 ACL on the interface.
 - h. Repeat the steps **b** to **g**.

201606140492

- Symptom: Traffic cannot be forwarded out of a 10-GE interface.
- Condition: This symptom occurs if the following tasks are performed:
 - a. Connect the 10-GE interface to an interface on the peer device.
 - b. Bind the 10-GE interface to an IRF port.
 - c. Remove the 10-GE interface from the IRF port.

201605040017

- Symptom: When the OSPF cost of an interface is modified, BFD session flapping occurs.
- Condition: This symptom might occur if the OSPF cost of an interface is modified when OSPF ECMP routes are load sharing traffic.

201606040084

- Symptom: On an IRF 3.1 system, the **display device** command displays information only about some PEXs.
- Condition: This symptom might occur if the parent fabric of an IRF 3.1 system is connected to 20 PEXs, and the cascade ports on the parent fabric are repeatedly shut down and brought up.

201606040039/201603110411

- Symptom: An IRF fabric does not respond to commands when a large number of ACLs exist.
- Condition: This symptom might occur if a large number of ACLs are configured on an IRF fabric.

201606020126/201605120177

- Symptom: After certain operations, the system prompts that resources are insufficient when a centralized VXLAN IP gateway group is specified on the switch.
- Condition: This symptom might occur if the operations are performed:
 - a. Execute the **vtep group member remote** command. VXLAN tunnels are automatically set up to the specified gateway group.
 - b. Execute the **undo vtep group member remote** command.
 - c. Restart the tunnel process.
 - d. Execute the **vtep group member remote** command.

201606010037/201606030468

- Symptom: When member switches in an IRF fabric are rebooted, the BGP process restarts unexpectedly.
- Condition: This symptom might occur if member switches in an IRF fabric are rebooted.

201605300311

- Symptom: When OpenFlow is enabled, an error is returned for the get event stream request.
- Condition: This symptom might occur if OpenFlow is enabled.

201605230206

- Symptom: In an EVPN network, the switch cannot receive EVPN IMET routes from a peer.
- Condition: This symptom might occur if EVPN is configured on the switch and its peer, and the two devices establish a BGP EVPN neighbor relationship.

201605190637

- Symptom: The ARP entry limit on an interface does not change after the table capacity mode is modified.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure an interface to send gratuitous ARP packets periodically.
 - b. Change the table capacity mode by using the **switch-mode** command, save the configuration, and reboot the switch.

201605130067

- Symptom: Two VMs access an EVPN network through the same VTEP. When multipath detection is performed for the VMs, the VTEP displays an incorrect outgoing interface for the first hop.
- Condition: This symptom might occur if multipath detection is performed for two VMs that access an EVPN network through the same VTEP.

201605100350

- Symptom: When the number of VSI interfaces exceeds the limit, the switch displays the incorrect notification "The parameter is incorrect."
- Condition: This symptom might occur if the **interface vsi-interface** command is repeatedly executed to create VSI interfaces, and the number of VSI interfaces exceeds the limit.

201605100029

- Symptom: When a VSI interface on a VXLAN IP gateway uses multiple IP addresses to provide gateway service, hosts cannot obtain the ARP information for the secondary IP addresses of the VSI interface.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure the switch as a VXLAN IP gateway, and assign multiple IP addresses to a VSI interface.
 - b. Repeatedly shut down and bring up the VSI interface.

201605070017

- Symptom: The help information for the **undo ip vpn-instance ?** command is incorrect.
- Condition: This symptom might occur if the following operations are performed:
 - a. Create a VPN instance.
 - b. Execute the **undo ip vpn-instance ?** command in system view.

201605030136

- Symptom: A merge operation by using invalid IfUnknownDomain data is successfully performed on the Domain/GlobalConfig table.
- Condition: This symptom might occur if a merge operation by using invalid IfUnknownDomain data is performed on the Domain/GlobalConfig table.

201604270260

- Symptom: A centralized VXLAN IP gateway group provides gateway service for a two-chassis IRF fabric that acts as a VTEP. After a master/subordinate switchover, the IRF fabric cannot send packets to all gateways because not all tunnels automatically established between the IRF fabric and the gateway group are assigned to the corresponding VXLAN.
- Condition: This symptom might occur if the following conditions exist:
 - A centralized VXLAN IP gateway group and an IRF fabric establish VXLAN tunnels automatically.
 - An IRF master/subordinate switchover occurs.

201606080442

- Symptom: LLDP cannot discover neighbors on an interface that uses an IPv6 address.
- Condition: This symptom might occur if an IPv6 address is assigned to an interface, and LLDP is enabled.

201604210596

- Symptom: A two-chassis IRF fabric that acts as a VXLAN IP gateway cannot forward Layer 3 traffic after a master/subordinate switchover.
- Condition: This symptom might occur if a master/subordinate switchover occurs in a two-chassis IRF fabric that acts as a VXLAN IP gateway.

201603310384

- Symptom: The system image is damaged if the switch is rebooted after it is successfully upgraded from the BootWare menus.
- Condition: This symptom might occur if the switch is rebooted after it is successfully upgraded from the BootWare menus.

201601260147

- Symptom: When local proxy ARP is enabled on distributed EVPN gateways, a gateway receives an ARP request for the ARP information of the gateway from a remote VM.
- Condition: This symptom might occur if local proxy ARP is enabled on distributed EVPN gateways.

201605070043

- Symptom: When an IRF fabric acts as a VTEP in an EVPN network, an IRF member cannot forward untagged traffic.
- Condition: This symptom might occur if a member in an IRF fabric receives untagged traffic.

201604110486

- Symptom: ACL resources are not released when relay entries on the DHCP relay agent are cleared.
- Condition: This symptom might occur if the following operations are performed:
 - a. Enable DHCP relay agent.
 - b. Execute the **reset dhcp relay client-information** command.

201607070519

- Symptom: After certain operations, a user-side aggregation group on two PEXs of the same PEX group cannot forward traffic.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure a user-side aggregation group on two PEXs by using the **interface bridge-aggregation interface-number pex** command.

- b. Assign one PEX to another PEX group, and then move it to the original PEX group.

201606130518

- Symptom: PEXs cannot come online when the parent fabric is repeatedly rebooted.
- Condition: This symptom might occur if link aggregation protocol flapping occurs on cascade ports of the parent fabric when the parent fabric is repeatedly rebooted.

201606030431

- Symptom: On the parent fabric of an IRF system, some TRILL access ports cannot forward broadcast traffic.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure multiple TRILL access ports on the parent fabric of an IRF system.
 - b. Disable TRILL globally.
 - c. Roll back to the configuration before TRILL is globally disabled.

201605310226

- Symptom: The parent fabric of an IRF system cannot forward the tagged traffic received from PEXs.
- Condition: This symptom might occur if tagged traffic received by PEXs is sent to a subordinate switch in the parent fabric.

201605240294

- Symptom: PEXs in an IRF system send traffic of VLAN 1 to the parent fabric. The traffic is flooded on TRILL access ports on IRF members except for the members that receive the traffic from the PEXs.
- Condition: This symptom might occur if PEXs send traffic of VLAN 1 to the parent fabric.

201605180288

- Symptom: In the output from the **display mac-address** command, the nicknames of egress RBs are not displayed.
- Condition: This symptom might occur if TRILL is enabled on an IRF 3.1 system.

201605270226

- Symptom: Specified port sends cc packet with rdi field set.
- Condition:
 - a. Configure cfd function.
 - b. Cfd cc packet without optional tlv fields is received on specified port.

201604120253

- Symptom: The switch cannot learn routes from two OSPF LSAs.
- Condition: This symptom might occur if two OSPF LSAs from a neighbor contain different information for the same transnet link.

201605120317/201605120311

- Symptom: The controller connected to the switch fails to issue flow entries and reply to ARP requests after certain operations are performed.
- Condition: This symptom occurs if the following conditions exist:
 - o An interface of the switch is configured with a large number of VLAN-VXLAN mappings.
 - o The switch receives unknown ARP packets.

Resolved problems in R2507

201606150562

- Symptom: When a QSFP-40G-LR4L-WDM1300 or QSFP-40G-ER4-WDM1300 transceiver module is installed in an interface, the speed of the interface becomes 100000 Mbps.
- Condition: This symptom might occur if a QSFP-40G-LR4L-WDM1300, QSFP-40G-ER4-WDM1300 or QSFP+ 40GBASE BIDI Optical Transceiver Module (850nm, 100m, SR) transceiver module is installed in an interface.

201606150192

- Symptom: Inter-device VXLAN connectivity is lost after certain operations.
- Condition: This symptom might occur if the following operations are performed:
 - a. Execute the **tunnel all** command in VXLAN view.
 - b. Use the **interface tunnel tunnel-number mode vxlan** command in system view to create a VXLAN tunnel interface.

201606130129

- Symptom: On the front panel of an 5940 48XGT 6QSFP28 /5940 48SFP+ 6QSFP28 switch, the LEDs of the interfaces numbered 49, 51, and 54 do not flash when the interfaces are forwarding traffic or 40-GE cables are installed in the interfaces.
- Condition: This symptom might occur if traffic is being forwarded by the interfaces numbered 49, 51, and 54 on the front panel of a 5940 48XGT 6QSFP28 /5940 48SFP+ 6QSFP28 switch or 40-GE cables are installed in the interfaces.

201606020318

- Symptom: The switch reboots multiple times after certain operations.
- Condition: This symptom might occur if the following operations are performed:
 - a. Configure the **classification vlan vlan-id loosen** command for an OpenFlow instance.
 - b. Save the configuration and reboot the switch.

201605240273

- Symptom: Speed negotiation fails on a 10-GE fiber port that connects to a 1-Gbps copper port of a peer device.
- Condition: This symptom might occur if a 10-GE fiber port installed with a 1000BASE-T 1G transceiver module is directly connected to a 1-Gbps copper port of a peer device.

TB201605030419

- Symptom: Routing policy configuration cannot be updated and hardware resources become insufficient after certain operations.
- Condition: This symptom might occur if the following operations are performed:
 - a. Apply a routing policy to a VXLAN tunnel interface to specify the next hop IP address.
 - b. Shut down and bring up the outgoing interface to the next hop IP address.

201605110274

- Symptom: After an IRF master/subordinate switchover, the laggd process exits unexpectedly during the reboot process of IRF members.
- Condition: This symptom might occur if an IRF master/subordinate switchover occurs,

201605100438

- Symptom: The switch cannot forward underlay traffic when ACs are mapped to VXLAN VSIs and then the VXLAN VSIs are deleted.

- Condition: This symptom might occur if ACs are mapped to VXLAN VSIs and then the VXLAN VSIs are deleted.

201605100250

- Symptom: When the switch acts as a VTEP and the VXLAN hardware resource allocation mode is set to Layer 2 gateway, an ACL cannot match incoming VXLAN packets by the inner IP header.
- Condition: This symptom might occur if the VXLAN hardware resource allocation mode is set to Layer 2 gateway, and the VXLAN tunnel source interface is a Layer 3 interface.

201605090001

- Symptom: As a VTEP, the switch cannot correctly forward double tagged frames that match an Ethernet service instance.
- Condition: This symptom might occur if the **encapsulation s-vid vlan-id c-vid vlan-id-list** command is configured for an Ethernet service instance.

201605060500

- Symptom: When certain operations are performed on an IRF fabric that acts as a VTEP, ECMP resources are exhausted, and VXLAN tunnels cannot come up.
- Condition: This symptom might occur if the following operations are performed:
 - a. Connect the IRF fabric to a gateway device through two Layer 3 aggregate interfaces.
 - b. Repeatedly shut down and bring up the Layer 3 aggregate interfaces on the gateway device.

TB201605040504

- Symptom: On an IRF fabric formed by 5940 48XGT 6QSFP+ and 5940 48SFP+ 6QSFP+ switches, the management Ethernet interface of the master cannot be pinged after a master/subordinate switchover.
- Condition: This symptom might occur if a master/subordinate switchover occurs in an IRF fabric formed by 5940 48XGT 6QSFP+ and 5940 48SFP+ 6QSFP+ switches.

201604210608

- Symptom: When all interfaces on the switch are shut down and then brought up, VXLAN tunnel configuration cannot be restored.
- Condition: This symptom might occur if all interfaces on the switch are shut down and then brought up.

201604220313

- Symptom: On a transport switch of a VXLAN network, an ACL cannot match VXLAN packets by the inner IP header.
- Condition: This symptom might occur if an ACL is configured to match VXLAN packets on a transport switch of a VXLAN network.

201603290139

- Symptom: The switch reboots unexpectedly when a Layer 3 Ethernet interface is split into four breakout interfaces.
- Condition: This symptom might occur if a Layer 3 Ethernet interface is split into four breakout interfaces while the interface is under attack of packets.

201512091309

- Symptom: The switch does not respond to commands in TCL configuration view when SSH login is used.

- Condition: This symptom might occur if SSH is used to log in to the switch and commands are executed in TCL configuration view.

201603250305

- Symptom: The switch fails to issue an ACL that contains the fragment parameter.
- Condition: This symptom might occur if the fragment parameter is specified for an ACL.

201603160814

- Symptom: On a 5940 IRF fabric configured with VXLANs, an aggregate interface sends the ARP packets received on a member port out of another member port.
- Condition: This symptom might occur if the following conditions exist:
 - a. VXLANs and a multichassis aggregate interface are configured on a 5940 IRF fabric.
 - b. Aggregation member ports have state changes when some IRF member switches are rebooted or some aggregation member ports are shut down.

201603090500

- Symptom: STP status of a port is not correct.
- Condition: This symptom occurs after the following operations are performed:
 - a. Create an aggregation group.
 - b. Enable or disable STP globally on the local device.
 - c. Bring up or shut down an aggregation member port in the aggregation group on the peer device.

201603030346

- Symptom: A user-defined queue scheduling profile uses byte-count WRR for a queue. After a reboot, weight-based WRR is used for the queue.
- Condition: This symptom might occur if the following operations are performed:
 - a. Create a queue scheduling profile, and configure byte-count WRR for a queue.
 - b. Delete the .mdb configuration file.
 - c. Save the running configuration and reboot the switch.

201604200046

- Symptom: In an IRF fabric with multidevice link aggregation, protocol flapping occurs on all link aggregation groups.
- Condition: This symptom occurs after the following operations are performed on an aggregation group:
 - a. Configure the aggregate interface as a trunk port and assign it to all VLANs by using the **port trunk permit vlan all** command.
 - b. Configure the aggregation group to operate in dynamic aggregation mode by using the **link-aggregation mode dynamic** command.
 - c. Configure the aggregation group to operate in static aggregation mode by using the **undo link-aggregation mode** command.
 - d. Configure the aggregation group to operate in dynamic aggregation mode by using the **link-aggregation mode dynamic** command.

201604250473

- Symptom: After an Ethernet service instance on an aggregate interface is deleted, the aggregate interface cannot forward traffic.
- Condition: This symptom occurs if the following operations are performed:
 - a. Create an Ethernet service instance on a Layer 2 aggregate interface.

- b. Use the **encapsulation untagged** command to configure the Ethernet service instance to match packets without VLAN tags.
- c. Use the **undo service-instance** command to delete the Ethernet service instance on the Layer 2 aggregate interface.

Resolved problems in R2506

First release.

Related documentation

The following documents provide related information:

- HPE FlexFabric 5940 Command References-Release 671x
- HPE FlexFabric 5940 Configuration Guides- Release 671x
- HPE FlexFabric 5940 Switch Series Installation Guide
- HPE LSWM124XG2Q (JH181A) & LSWM124XG2QL (JH180A) Interface Modules User Guide
- HPE LSWM124XGT2Q Interface Module (JH182A) User Guide
- HPE LSWM18QC Interface Module (JH183A) User Guide
- HPE LSWM12H2Q Interface Module (JH409A) User Guide
- HPE A58x0AF 650W AC (JC680A) & 650W DC(JC681A) Power Supplies User Guide
- HPE FlexFabric Switch 650W 48V Hot Plug NEBS Compliant DC Power Supply 650W DC (JH336A) User Guide
- HPE A58x0AF 300W AC (JG900A) & 300W DC(JG901A) Power Supplies User Guide
- HPE LSWM1HFANSC and LSWM1HFANSCB Fan Assemblies Installation
- HPE LSWM1FANSC & LSWM1FANSCB Fan Trays User Guide
- HPE LSWM1BFANSC & LSWM1BFANSCB Fan Trays User Guide

Support and other resources

Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:
www.hpe.com/assistance
- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:
www.hpe.com/support/hpesc

Information to collect:

- Technical support registration number (if applicable).
- Product name, model or version, and serial number.
- Operating system name and version.
- Firmware version.
- Error messages.
- Product-specific reports and logs.
- Add-on products or components.

- Third-party products or components.

Documents

To find related documentation, see the Hewlett Packard Enterprise Support Center website at <http://www.hpe.com/support/hpesc>.

- Enter your product name or number and click **Go**. If necessary, select your product from the resulting list.
- For a complete list of acronyms and their definitions, see HPE FlexNetwork technology acronyms.

Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (docsfeedback@hpe.com). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.

Appendix A Feature list

Hardware features

Table 5 5940 series hardware features

| Item | FlexFabric 5940 32QSFP+ JH396A | FlexFabric 5940 48SFP+ 6QSFP+ JH395A | FlexFabric 5940 48XGT 6QSFP+ JH394A | FlexFabric 5940 48SFP+ 6QSFP28 JH390A | FlexFabric 5940 48XGT 6QSFP28 JH391A |
|---------------------------|---|---|---|---|---|
| Dimensions (H × W × D) | 43.6 × 440 × 660 mm (1.72 × 17.32 × 25.98 in) | 43.6 × 440 × 460 mm (1.72 × 17.32 × 18.11 in) | 43.6 × 440 × 660 mm (1.72 × 17.32 × 25.98 in) | 43.6 × 440 × 460 mm (1.72 × 17.32 × 18.11 in) | 43.6 × 440 × 660 mm (1.72 × 17.32 × 25.98 in) |
| Weight | ≤ 13.0 kg (28.66 lb) | ≤ 10kg (22.04lb) | ≤ 13.0 kg (28.66 lb) | ≤ 10kg (22.04lb) | ≤ 13.0 kg (28.66 lb) |
| Console ports | 1 | 1 | 1 | 1 | 1 |
| Mini USB(Console) | 1 | N/A | N/A | 1 | 1 |
| Management Ethernet ports | 1 | 1 | 1 | 2 | 2 |
| USB ports | 1 | | | | |
| 10G Base-T Ethernet ports | N/A | N/A | 48 | N/A | 48 |
| SFP+ ports | N/A | 48 | N/A | 48 | N/A |
| QSFP+ ports | 32 | 6 | 6 | N/A | N/A |
| Fan trays | 2 | 2 | 2 | 2 | 2 |
| interface card model | N/A | N/A | N/A | N/A | N/A |
| Power modules | 2 | 2 | 2 | 2 | 2 |
| AC-input voltage | Rated voltage: 100 VAC to 240 V AC @ 50 or 60 Hz Max voltage: 90 VAC to 264 V AC @ 47 to 63 Hz | | | | |
| DC-input voltage | Rated voltage: -40 VDC to -60 VDC Max voltage: -40 VDC to -72 VDC | | | | |

| Item | FlexFabric 5940 32QSFP+ JH396A | FlexFabric 5940 48SFP+ 6QSFP+ JH395A | FlexFabric 5940 48XGT 6QSFP+ JH394A | FlexFabric 5940 48SFP+ 6QSFP28 JH390A | FlexFabric 5940 48XGT 6QSFP28 JH391A |
|---------------------------|--|--|--|--|--|
| Minimum power consumption | Single-AC: 132W Dual-AC: 145W Single-DC: 128W Dual-AC: 142W | Single AC input: 75W Dual-DC: 81W Single AC input: 72W Dual AC inputs: 80W | Single DC input: 135W Dual DC inputs: 150W Single DC input: 132W Dual AC inputs: 142W | Single DC input: 89W Dual DC inputs: 95W Single DC input: 91W Dual AC inputs: 99W | Single DC input: 142W Dual DC inputs: 157W Single DC input: 135W Dual AC inputs: 150W |
| Maximum power consumption | Single-AC: 291W Dual-AC: 301W Single-DC: 291W Dual-DC: 299W | Single AC input: 145W Dual AC inputs: 150W Single DC input: 141W Dual DC inputs: 145W | Single-AC: 260W Dual-AC: 270W Single-DC: 251W Dual-DC: 265W | Single-AC: 186W Dual-AC: 196W Single-DC: 187W Dual-DC: 192W | Single-AC: 310W Dual-AC: 320W Single-DC: 294W Dual-DC: 301W |
| Operating temperature | 0°C to 45°C (32°F to 113°F) | | | | |
| Operating humidity | 10% to 90%, noncondensing | | | | |

Table 6 FF 5940 2-slot Switch/5940 4-slot Switch series hardware features

| Item | FF 5940 2-slot Switch JH397A | FF 5940 4-slot Switch JH398A |
|---------------------------|--|--|
| Dimensions (H × W × D) | 44.2 × 440 × 660 mm (1.74 × 17.32 × 25.98 in) | 88.1 × 440 × 660 mm (3.46 × 17.32 × 25.98 in) |
| Weight | ≤16kg (35.27lb) | ≤27kg (59.52lb) |
| Console ports | 1 | 1 |
| Mini USB(Console) | 1 | N/A |
| Management Ethernet ports | 1 | 1 |
| USB ports | 1 | |
| 10G Base-T Ethernet ports | N/A | N/A |
| SFP+ ports | N/A | N/A |

| Item | FF 5940 2-slot Switch JH397A | FF 5940 4-slot Switch JH398A |
|--------------------------------------|--|--|
| QSFP+ ports | 2 | N/A |
| Fan trays | LSVM1HFANSC LSVM1HFANSCB | LSWM1BFANSC LSWM1BFANSCB |
| interface card model | <ul style="list-style-type: none"> • HPE 5940 2-port QSFP+ and 2-port QSFP28 Module JH409A • HPE 5930 8-port QSFP+ Module JH183A • HPE 5930 24p SFP+ and 2p QSFP+ Mod JH180A • HPE 5930 24p SFP+ and 2p QSFP+ Mod w Msec JH181A • HPE 5930 24p 10GBase-T and 2p QSFP+ Mod JH182A • HPE 5930 24p CP and 2p QSFP+ Mod JH184A • HPE FlexFabric 5950 8-port QSFP28 MACsec Module JH957A | |
| Power modules | 2 | 2 |
| AC-input voltage | <ul style="list-style-type: none"> • Rated voltage: 100 VAC to 240 VAC @ 50 or 60 Hz • Max voltage: 90 VAC to 264 VAC @ 47 to 63 Hz | |
| DC-input voltage | <ul style="list-style-type: none"> • Rated voltage: –40 VDC to –60 VDC • Max voltage: –40 VDC to –72 VDC | |
| Minimum power consumption | <ul style="list-style-type: none"> • Single AC input: 95 W • Dual AC inputs: 110 W • Single DC input: 91 W • Dual DC inputs: 105 W | <ul style="list-style-type: none"> • Dual AC inputs: 135 W • Triple AC inputs: 150 W • Quadruple AC inputs: 165 W • Dual DC inputs: 131 W • Triple DC inputs: 145 W • Quadruple DC inputs: 155 W |
| Maximum power consumption | <ul style="list-style-type: none"> • Single AC input: 439 W • Dual AC inputs: 450 W • Single DC input: 443 W • Dual DC inputs: 445 W | <ul style="list-style-type: none"> • Dual AC inputs: 827 W • Triple AC inputs: 837 W • Quadruple AC inputs: 856 W • Dual DC inputs: 823 W • Triple DC inputs: 825 W • Quadruple DC inputs: 828 W |
| Melting current of power module fuse | <ul style="list-style-type: none"> • LSVM1AC650: 10 A @ 250 VAC • LSVM1DC650: 30 A @ 250 VDC | |
| Operating temperature | 0°C to 45°C (32°F to 113°F) | |
| Operating humidity | 10% to 90%, noncondensing | |
| Fire resistance compliance | UL60950-1, EN60950-1, IEC60950-1, GB4943 | |

Software features

Table 7 Software features of the 5940 series

| Feature | FlexFabric 5940 32QSFP+ JH396A | FlexFabric 5940 48SFP+ 6QSFP+ JH395A | FlexFabric 5940 48XGT 6QSFP+ JH394A | FlexFabric 5940 48SFP+ 6QSFP28 JH390A | FlexFabric 5940 48XGT 6QSFP28 JH391A |
|---|---|--------------------------------------|-------------------------------------|---------------------------------------|--------------------------------------|
| Full duplex Wire speed L2 switching capacity | 2560Gbps | 2160Gbps | 2160Gbps | 1440Gbps | 1440Gbps |
| Whole system Wire speed L2 switching Packet forwarding rate | 1904.64 | 1607.04 | 1607.04 | 1011.84 | 1011.84 |
| Forwarding mode | Store-forward and cut-through | | | | |
| IRF | <ul style="list-style-type: none"> • Ring topology • Daisy chain topology • LACP MAD • ARP MAD • ND MAD • BFD MAD • ISSU • IRF comprised of different models | | | | |
| Link aggregation | <ul style="list-style-type: none"> • Aggregation of 10-GE ports • Aggregation of 40-GE ports • Static link aggregation • Dynamic link aggregation • When stacked, supports up to 1024 aggregation groups, each supporting up to 32 ports | | | | |
| Flow control | IEEE 802.3x flow control and back pressure | | | | |
| Data center | <ul style="list-style-type: none"> • PFC • DCBX • FCoE(FCF/Transit/NPV) • EVB • TRILL • OpenFlow • VxLAN • VxLAN L3 Gateway • EVPN | | | | |
| Jumbo Frame | Supports maximum frame size of 10000 | | | | |
| MAC address table | <ul style="list-style-type: none"> • 288K MAC addresses • 1K static MAC addresses • Blackhole MAC addresses • MAC address learning limit on a port | | | | |

| | |
|------------------------|--|
| VLAN | <ul style="list-style-type: none"> • Port-based VLANs (4094 VLANs) • QinQ and selective QinQ |
| VLAN mapping | <ul style="list-style-type: none"> • One-to-one VLAN mapping • Many-to-one VLAN mapping • Two-to-two VLAN mapping |
| ARP | <ul style="list-style-type: none"> • 128K entries • 1K static entries • Gratuitous ARP • Standard proxy ARP and local proxy ARP • ARP source suppression • ARP black hole • ARP detection (based on DHCP snooping entries/802.1x security entries/static IP-to-MAC bindings) • Multicast ARP • ARP logging • IRDP • ARP proxy |
| ND | <ul style="list-style-type: none"> • 64K entries • 1K static entries |
| VLAN virtual interface | 2K |
| DHCP | <ul style="list-style-type: none"> • DHCP client • DHCP snooping • DHCP relay agent • DHCP server • DHCPv6 server |
| UDP helper | Supported |
| DNS | <ul style="list-style-type: none"> • Dynamic domain name resolution • Dynamic domain name resolution client • IPv4/IPv6 addresses |
| IPv4 routing | <ul style="list-style-type: none"> • 4K static routes • RIP (Routing Information Protocol) v1/v2; up to 4K IPv4 routes • OSPF (Open Shortest Path First) v1/v2; up to 250K IPv4 routes • BGP (Border Gateway Protocol); up to 250K IPv4 routes • ISIS (Intermediate System to Intermediate system; up to 250K IPv4 routes • Configurable maximum number of equal-cost routes; up to 16K equal-cost routes • VRRP • PBR |
| IPv6 routing | <ul style="list-style-type: none"> • 2K static routes • RIPng: Supports up to 2K IPv6 routes • OSPF v3: Supports up to 250K IPv6 routes • ISISv6: Supports up to 250K IPv6 routes • Up to 16K ECMP routes; each ECMP route supports up to 128 next hops • Routing policy • VRRP • PBR |
| URPF | Reverse route check strict mode and loose mode |
| MCE | Supported |

| | |
|---|---|
| BFD | <ul style="list-style-type: none"> • OSPF/OSPFv3 • BGP/BGP4 • IS-IS/IS-ISv6 • PIM/IPM for IPv6 • Static route • MAD |
| Tunnel | <ul style="list-style-type: none"> • IPv4 over IPv4 tunnel • IPv4 over IPv6 tunnel • IPv6 over IPv4 manual tunnel • IPv6 over IPv4 6to4 tunnel • IPv6 over IPv4 ISATAP tunnel • IPv6 over IPv6 tunnel • GRE tunnel |
| IPv4 multicast | <ul style="list-style-type: none"> • IGMP snooping v1/v2/v3 • IGMP v1/v2/v3 • PIM-DM • PIM-SM |
| IPv6 multicast | <ul style="list-style-type: none"> • MLD snooping v1/v2 • MLD v1/v2 • PIM-DM/SM for IPv6 |
| Broadcast/multicast/unicast storm control | <ul style="list-style-type: none"> • Storm control based on port rate percentage • PPS-based storm control • Bps-based storm control |
| MSTP | <ul style="list-style-type: none"> • STP/RSTP/MSTP protocol • STP Root Guard • BPDU Guard |
| QoS/ACL | <ul style="list-style-type: none"> • Restriction of the rates at which a port sends and receives packets, with a granularity of 8 kbps. • Packet redirect • Committed access rate (CAR), with a granularity of traffic limit 8 kbps. • Eight output queues for each port • Flexible queue scheduling algorithms based on port and queue, including strict priority (SP), Weighted Deficit Round Robin (WDRR), Weighted Fair Queuing (WFQ), SP + WDRR, and SP + WFQ. • Remarking of 802.1p and DSCP priorities • Packet filtering at L2 (Layer 2) through L4 (Layer 4); flow classification based on source MAC address, destination MAC address, source IP (IPv4/IPv6) address, destination IP (IPv4/IPv6) address, port, protocol, and VLAN. • Time range • Weighted Random Early Detection (WRED) • Queue shaping • User profile • COPP • Explicit Congestion Notification (ECN) |
| Mirroring | <ul style="list-style-type: none"> • Flow mirroring • Port mirroring • Multiple mirror observing port |
| Remote mirroring | <ul style="list-style-type: none"> • Port remote mirroring (RSPAN) • ERSPAN |

| | |
|-----------------------|---|
| Security | <ul style="list-style-type: none"> • Hierarchical management and password protection of users • AAA authentication • RADIUS authentication • HWTACACS • SSH 2.0 • Port isolation • Port security • IP-MAC-port binding • IP Source Guard • HTTPS • SSL • PKI • Boot ROM access control (password recovery) |
| 802.1X | <ul style="list-style-type: none"> • Up to 2,048 users • Port-based and MAC address-based authentication • Trunk port authentication |
| Traffic Management | <ul style="list-style-type: none"> • sFlow |
| Loading and upgrading | <ul style="list-style-type: none"> • Loading and upgrading through XModem protocol • Loading and upgrading through FTP • Loading and upgrading through the trivial file transfer protocol (TFTP) |
| Management | <ul style="list-style-type: none"> • Configuration at the command line interface • Remote configuration through Telnet • Configuration through Console port • Python • NETCONF • Simple network management protocol (SNMP) • Neutron, Chef, Puppet, RESTful • System log • Hierarchical alarms • NTP • Power supply alarm function • Fan and temperature alarms |
| Maintenance | <ul style="list-style-type: none"> • Debugging information output • Ping and Tracert • NQA • Track • Remote maintenance through Telnet • 802.1ag • 802.3ah • DLDP • File download and upload through USB port |
| M-LAG | <ul style="list-style-type: none"> • Support NETCONF • Traffic forwarding • Support Multicast • MVPN • Multicast base VXLAN |

Appendix B Fixed security vulnerabilities

Fixed security vulnerabilities in R6710

CVE-2020-7469

In FreeBSD 12.2-STABLE before r367402, 11.4-STABLE before r368202, 12.2-RELEASE before p1, 12.1-RELEASE before p11 and 11.4-RELEASE before p5 the handler for a routing option caches a pointer into the packet buffer holding the ICMPv6 message. However, when processing subsequent options the packet buffer may be freed, rendering the cached pointer invalid. The network stack may later dereference the pointer, potentially triggering a use-after-free.

CVE-2021-22924

libcurl keeps previously used connections in a connection pool for subsequent transfers to reuse, if one of them matches the setup. Due to errors in the logic, the config matching function did not take 'issuercert' into account and it compared the involved paths *case insensitively*, which could lead to libcurl reusing wrong connections. File paths are, or can be, case sensitive on many systems but not all, and can even vary depending on used file systems. The comparison also didn't include the 'issuer cert' which a transfer can set to qualify how to verify the server certificate.

Fixed security vulnerabilities in F6705

CVE-2022-0778

The BN_mod_sqrt() function, which computes a modular square root, contains a bug that can cause it to loop forever for non-prime moduli. Internally this function is used when parsing certificates that contain elliptic curve public keys in compressed form or explicit elliptic curve parameters with a base point encoded in compressed form. It is possible to trigger the infinite loop by crafting a certificate that has invalid explicit curve parameters. Since certificate parsing happens prior to verification of the certificate signature, any process that parses an externally supplied certificate may thus be subject to a denial of service attack. The infinite loop can also be reached when parsing crafted private keys as they can contain explicit elliptic curve parameters. Thus vulnerable situations include: - TLS clients consuming server certificates - TLS servers consuming client certificates - Hosting providers taking certificates or private keys from customers - Certificate authorities parsing certification requests from subscribers - Anything else which parses ASN.1 elliptic curve parameters. Also any other applications that use the BN_mod_sqrt() where the attacker can control the parameter values are vulnerable to this DoS issue. In the OpenSSL 1.0.2 version the public key is not parsed during initial parsing of the certificate which makes it slightly harder to trigger the infinite loop. However any operation which requires the public key from the certificate will trigger the infinite loop. In particular the attacker can use a self-signed certificate to trigger the loop during verification of the certificate signature. This issue affects OpenSSL versions 1.0.2, 1.1.1 and 3.0. It was addressed in the releases of 1.1.1n and 3.0.2 on the 15th March 2022. Fixed in OpenSSL 3.0.2 (Affected 3.0.0, 3.0.1). Fixed in OpenSSL 1.1.1n (Affected 1.1.1-1.1.1m). Fixed in OpenSSL 1.0.2zd (Affected 1.0.2-1.0.2zc).

Fixed security vulnerabilities in R2702

CVE-2018-5407

OpenSSL is prone to a local information-disclosure vulnerability. Local attackers can exploit this issue to obtain sensitive information. This may aid in further attacks.

CVE-2018-15473

OpenSSH is prone to a user-enumeration vulnerability. An attacker may leverage this issue to harvest valid user accounts, which may aid in brute-force attacks. OpenSSH through 7.7 are vulnerable; other versions may also be affected.

Fixed security vulnerabilities in R2612

CVE-2017-12190

Local attacker can exploit these issues to obtain sensitive information that may lead to further attacks.

CVE-2017-12192

Attackers can exploit this issue to cause denial-of-service conditions. Due to the nature of this issue, arbitrary code execution may be possible but this has not been confirmed.

CVE-2017-15274

An attacker can exploit this issue to cause a local denial-of-service condition.

CVE-2017-15299

An attacker can exploit this issue to trigger a kernel panic, denying service to legitimate users.

CVE-2017-3736

An attacker can exploit this issue to gain access to sensitive information that may aid in further attacks.

CVE-2017-15896

An attacker can exploit this issue to bypass TLS validate and encrypt, send application data to Node.js.

CVE-2017-3737

Successfully exploiting this issue will allow attackers to bypass security restrictions and perform unauthorized actions; this may aid in launching further attacks.

CVE-2017-3738

An attacker can exploit this issue to gain access to sensitive information that may aid in further attacks.

CVE-2017-1000253

Local attackers may exploit this issue to gain root privileges.

Fixed security vulnerabilities in R2610

CVE-2017-3735

Successfully exploiting this issue will allow attackers to bypass security restrictions and perform unauthorized actions; this may aid in launching further attacks.

Fixed security vulnerabilities in R2609

CVE-2017-6458

NTP are prone to a buffer-overflow vulnerability because it fails to properly bounds-check user-supplied data before copying it into an insufficiently sized buffer.

CVE-2016-9042

NTP is prone to a denial-of-service vulnerability. An attacker can exploit this issue to cause a denial-of-service condition, denying service to legitimate users.

Fixed security vulnerabilities in R2608P02

CVE-2016-6210

An attacker may leverage this issue to harvest valid user accounts, which may aid in brute-force attacks.

Fixed security vulnerabilities in F2608

CVE-2017-3731

OpenSSL is prone to denial-of-service vulnerability. An attacker may exploit this issue to crash the application, resulting in denial-of-service condition.

CVE-2017-3732

OpenSSL is prone to an information-disclosure vulnerability. An attacker can exploit this issue to gain access to sensitive information that may aid in further attacks.

CVE-2015-3405

An attacker can exploit the ntp-keygen utility to spoof an NTP client or server.

CVE-2014-9297

An attacker can exploit this issue. When an NTP client decrypted a secret received from an NTP server.

CVE-2015-9298

An attacker could bypass source IP restrictions and send malicious control and configuration packets.

Fixed security vulnerabilities in F2605

CVE-2016-7431:

Zero Origin timestamp problems were fixed by Bug 2945 in ntp-4.2.8p6. However, subsequent timestamp validation checks introduced a regression in the handling of some Zero origin timestamp checks.

CVE-2016-7428:

The broadcast mode of NTP is expected to only be used in a trusted network. If the broadcast network is accessible to an attacker, a potentially exploitable denial of service vulnerability in ntpd's broadcast mode poll interval enforcement functionality can be abused. To limit abuse, ntpd restricts the rate at which each broadcast association will process incoming packets. ntpd will reject broadcast mode packets that arrive before the poll interval specified in the preceding broadcast packet expires. An attacker with access to the NTP broadcast domain can send specially crafted broadcast mode NTP packets to the broadcast domain which, while being logged by ntpd, will cause ntpd to reject broadcast mode packets from legitimate NTP broadcast servers.

CVE-2016-7427:

The broadcast mode of NTP is expected to only be used in a trusted network. If the broadcast network is accessible to an attacker, a potentially exploitable denial of service vulnerability in ntpd's broadcast mode replay prevention functionality can be abused. An attacker with access to the NTP

broadcast domain can periodically inject specially crafted broadcast mode NTP packets into the broadcast domain which, while being logged by ntpd, can cause ntpd to reject broadcast mode packets from legitimate NTP broadcast servers.

Fixed security vulnerabilities in F2604

CVE-2016-6304:

Multiple memory leaks in t1_lib.c in OpenSSL before 1.0.1u, 1.0.2 before 1.0.2i, and 1.1.0 before 1.1.0a allow remote attackers to cause a denial of service (memory consumption) via large OCSP Status Request extensions.

CVE-2016-6306

The certificate parser in OpenSSL before 1.0.1u and 1.0.2 before 1.0.2i might allow remote attackers to cause a denial of service (out-of-bounds read) via crafted certificate operations, related to s3_clnt.c and s3_srvr.c.

Fixed security vulnerabilities in F2603

CVE-2016-8858

A remote user can send specially crafted data during the key exchange process to trigger a flaw in kex_input_kexinit() and consume excessive memory on the target system. This can be exploited to consume up to 384 MB per connection.

CVE-2016-5195

An unprivileged local user could use this flaw to gain write access to otherwise read-only memory mappings and thus increase their privileges on the system.

CVE-2009-3238

The get_random_int function in the Linux kernel before 2.6.30 produces insufficiently random numbers, which allows attackers to predict the return value, and possibly defeat protection mechanisms.

CVE-2016-0705

Fixed vulnerability when OpenSSL parses malformed DSA private keys and could lead to a DoS attack or memory corruption for applications that receive DSA private keys from untrusted sources.

CVE-2016-0798

Fixed vulnerability in OpenSSL 1.0.1 before 1.0.1s and 1.0.2 before 1.0.2g allows remote attackers to cause a denial of service (memory consumption) by providing an invalid username in a connection attempt.

CVE-2016-0797

Fixed vulnerability in OpenSSL 1.0.1 before 1.0.1s and 1.0.2 before 1.0.2g allow remote attackers to cause a denial of service (heap memory corruption or NULL pointer dereference).

CVE-2016-0799

Fixed vulnerability in OpenSSL 1.0.1 before 1.0.1s and 1.0.2 before 1.0.2g improperly calculates string lengths, which allows remote attackers to cause a denial of service which could lead to memory allocation failure or memory leaks.

CVE-2016-0702

Fixed vulnerability in OpenSSL 1.0.1 before 1.0.1s and 1.0.2 before 1.0.2g which makes it easier for local users to discover RSA keys leveraging cache-bank conflicts, aka a "CacheBleed" attack.

Fixed security vulnerabilities in R2509P02

CVE-2016-2177

Fixed vulnerability in `s3_srvr.c`, `ssl_sess.c`, and `t1_lib.c` functions in OpenSSL through 1.0.2h that allows remote attackers to cause a denial of service (integer overflow and application crash), or possibly have an unspecified other impact by leveraging unexpected malloc behavior.

Fixed security vulnerabilities in R2509

CVE-2016-4953

Fixed vulnerability in NTP 4.x before 4.2.8p8 allows remote attackers to cause a denial of service by sending a spoofed packet with incorrect authentication data at a certain time.

CVE-2016-4954

Fixed vulnerability in `ntpd` in NTP 4.x before 4.2.8p8 allows remote attackers to cause a denial of service by sending spoofed packets from source IP addresses in a certain scenario.

CVE-2016-4956

Fixed vulnerability in NTP 4.x before 4.2.8p8 allows remote attackers to cause a denial of service via a spoofed broadcast packet.

CVE-2016-2105

Fixed vulnerability in “EVP Encode” in OpenSSL before 1.0.1t and 1.0.2 before 1.0.2h allows remote attackers to cause a denial of service (heap memory corruption) via a large amount of binary data.

CVE-2016-2106

Fixed vulnerability in “EVP Encrypt” in OpenSSL before 1.0.1t and 1.0.2 before 1.0.2h allows remote attackers to cause a denial of service (heap memory corruption) via a large amount of binary data.

CVE-2016-2107

Fixed vulnerability in OpenSSL before 1.0.1t and 1.0.2h allows remote attackers to obtain sensitive cleartext information via a padding-oracle attack against an AES CBC session.

CVE-2016-2108

Fixed vulnerability in OpenSSL before 1.0.1o and 1.0.2 before 1.0.2c allows remote attackers to execute arbitrary code or cause a denial of service (buffer underflow and memory corruption).

CVE-2016-2109

Fixed vulnerability in “asn” before 1.0.1t and 1.0.2 before 1.0.2h allows remote attackers to cause a denial of service (memory consumption) via a short invalid encoding.

CVE-2016-2176

Fixed vulnerability in “X509” in OpenSSL before 1.0.1t and 1.0.2 before 1.0.2h allows remote attackers to obtain sensitive information from memory or cause a denial of service

Fixed security vulnerabilities in R2508

CVE-2016-1547

Fixed vulnerability where an off-path attacker can deny service to `ntpd` clients by demobilizing preemptable associations using spoofed crypto-NAK packets.

CVE-2016-1548

Fixed vulnerability where an attacker can change the time of an ntpd client or deny service to an ntpd client by forcing it to change from basic client/server mode to interleaved symmetric mode.

CVE-2016-1550

Fixed vulnerability in ntpd function allow an attacker to conduct a timing attack to compute the value of the valid authentication digest causing forged packets to be accepted by ntpd.

CVE-2016-1551

Fixed vulnerability in ntpd allows unauthenticated network attackers to spoof refclock packets to ntpd processes on systems that do not implement bogon filtering.

CVE-2016-2519

Fixed vulnerability in ntpd will abort if an attempt is made to read an oversized value.

CVE-2015-7704

Fixed vulnerability in ntpd that a remote attacker could use, to send a packet to an ntpd client that would increase the client's polling interval value, and effectively disable synchronization with the server.

CVE-2015-8138

Fixed vulnerability in ntpd which attackers may be able to disable time synchronization by sending a crafted NTP packet to the NTP client.

CVE-2015-7979

Fixed vulnerability in ntpd allows attackers to send special crafted broadcast packets to broadcast clients, which may cause the affected NTP clients to become out of sync over a longer period of time.

CVE-2015-7974

Fixed vulnerability in NTP 4.x before 4.2.8p6 and 4.3.x before 4.3.90 which might allow remote attackers to conduct impersonation attacks via an arbitrary trusted key.

CVE-2015-7973

Fixed vulnerability when NTP is configured in broadcast mode, a man-in-the-middle attacker or a malicious client could replay packets received from the broadcast server to all (other) clients, which cause the time on affected clients to become out of sync over a longer period of time.

Fixed security vulnerabilities in R2507

CVE-2016-0800

Fixed vulnerability where an attacker will decrypt TLS sessions by using a server supporting SSLv2 and EXPORT cipher suites.

CVE-2016-0705

Fixed vulnerability when OpenSSL parses malformed DSA private keys and could lead to a DoS attack or memory corruption for applications that receive DSA private keys from untrusted sources.

CVE-2016-0798

Fixed vulnerability in OpenSSL 1.0.1 before 1.0.1s and 1.0.2 before 1.0.2g allows remote attackers to cause a denial of service (memory consumption) by providing an invalid username in a connection attempt.

CVE-2016-0797

Fixed vulnerability in OpenSSL 1.0.1 before 1.0.1s and 1.0.2 before 1.0.2g allow remote attackers to cause a denial of service (heap memory corruption or NULL pointer dereference).

CVE-2016-0799

Fixed vulnerability in OpenSSL 1.0.1 before 1.0.1s and 1.0.2 before 1.0.2g improperly calculates string lengths, which allows remote attackers to cause a denial of service which could lead to memory allocation failure or memory leaks.

CVE-2016-0702

Fixed vulnerability in OpenSSL 1.0.1 before 1.0.1s and 1.0.2 before 1.0.2g which makes it easier for local users to discover RSA keys leveraging cache-bank conflicts, aka a "CacheBleed" attack.

CVE-2016-0703

Fixed vulnerability if an eavesdropper has intercepted an SSLv2 handshake, they can use the server as an oracle to determine the SSLv2 master-key.

CVE-2016-0704

Fixed vulnerability in the SSLv2 implementation in OpenSSL before 0.9.8zf, 1.0.0 before 1.0.0r, 1.0.1 before 1.0.1m, and 1.0.2 before 1.0.2a overwrites incorrect MASTER-KEY bytes during use of export cipher suites, which makes it easier for remote attackers to decrypt TLS ciphertext data.

CVE-2016-2842

Fixed vulnerability in the `doapr_outch` function in `crypto/bio/b_print.c`, which allows remote attackers to cause a denial of service (out-of-bounds write or memory consumption) or possibly have unspecified other impact via a long string.

Appendix C Upgrading software

This section describes how to upgrade system software while the router is operating normally or when the router cannot correctly start up.

System software file types

Software required for starting up the switch includes:

- **Boot ROM image**—A .bin file that comprises a basic section and an extended section. The basic section is the minimum code that bootstraps the system. The extended section enables hardware initialization and provides system management menus. You can use these menus to load software and the startup configuration file or manage files when the switch cannot correctly start up.
- **Software images**—Includes boot images and system images.
 - **Boot image**—A .bin file that contains the operating system kernel. It provides process management, memory management, file system management, and the emergency shell.
 - **System image**—A .bin file that contains the main application code required for device operation. This includes device management, interface management, configuration management, and routing management.

The software images that have been loaded are called "current software images." The software images specified to load at next startup are called "startup software images."

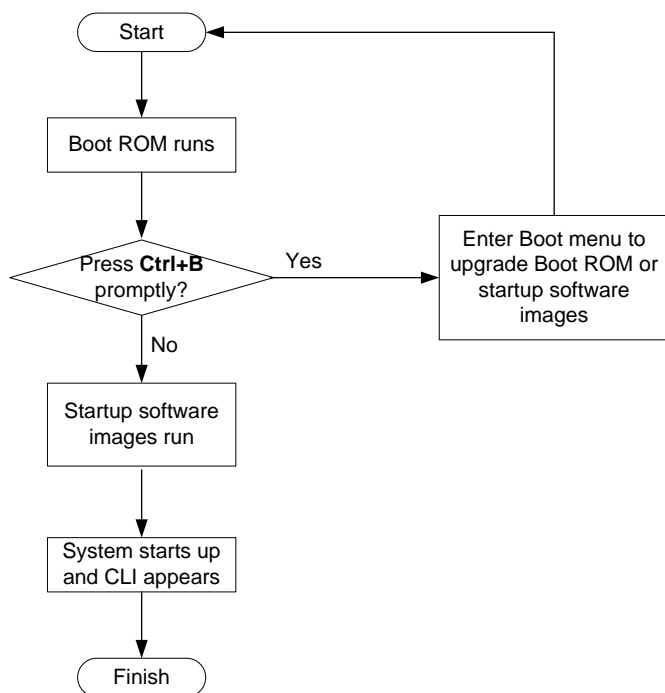
These images might be released separately or as a whole in one .ipe package file. If an .ipe file is used, the system automatically decompresses the file, loads the .bin boot and system images in the file and sets them as startup software images. Typically, the Boot ROM and software images for this switch series are released in an .ipe file named **main.ipe**.

In addition to these images, HPE irregularly releases patch packages for you to fix bugs without rebooting the switch. A patch package does not add new features or functions.

System startup process

Upon power-on, the Boot ROM image runs to initialize hardware and then the software images run to start up the entire system, as shown in [Figure 1](#).

Figure 1 System startup process



Upgrade methods

You can upgrade system software by using one of the following methods:

| Upgrading method | Software types | Remarks |
|------------------------|-----------------|---|
| Upgrading from the CLI | Software images | <ul style="list-style-type: none">You must reboot the switch to complete the upgrade.This method can interrupt ongoing network services. |
| | Patch packages | <p>The upgrade does not interrupt ongoing services.</p> <p>Make sure the patch packages match the current software images. A patch package can fix bugs only for its matching software image version.</p> |

| Upgrading method | Software types | Remarks |
|------------------------------|---|---|
| Upgrading from the Boot menu | <ul style="list-style-type: none"> • Boot ROM image • Software images | <p>Use this method when the switch cannot correctly start up.</p> <p>⚠ CAUTION:</p> <p>Upgrading an IRF fabric from the CLI rather than the Boot menu.</p> <p>The Boot menu approach requires that you upgrade the member switches one by one and has larger impact on services than the CLI approach.</p> |

The output in this document is for illustration only and might vary with software releases. For example, this document uses boot.bin and system.bin to represent boot and system image names, whereas the actual software image name format is chassis_software platform version_image type_release, for example, 5940-cmw710-boot-r2506.bin and 5940-cmw710-system-r2506.bin.

Upgrading from the CLI

This section uses a two-member IRF fabric as an example to describe how to upgrade software from the CLI. If you have more than two subordinate switches, repeat the steps for the subordinate switch to upgrade their software. If you are upgrading a standalone switch, ignore the steps for upgrading the subordinate switch.

Preparing for the upgrade

Before you upgrade software, complete the following tasks:

1. Log in to the IRF fabric through Telnet or the console port (details not shown).
2. Perform the **display irf** command in any view to identify the number of IRF members, each member switch's role and IRF member ID.

```
<Sysname> display irf
```

| MemberID | Role | Priority | CPU-Mac | Description |
|----------|---------|----------|----------------|-------------|
| *+1 | Master | 5 | 0023-8927-afdc | --- |
| 2 | Standby | 1 | 0023-8927-af43 | --- |

```
-----
* indicates the device is the master.
+ indicates the device through which the user logs in.
```

```
The Bridge MAC of the IRF is: 0023-8927-afdb
Auto upgrade           : no
Mac persistent         : 6 min
Domain ID              : 0
```

3. Perform the **dir** command in user view to identify the free storage space of each member switch.
4. Identify the free Flash space of the master switch.

```
<Sysname> dir
Directory of flash:
 0      -rw-      41424  Aug 23 2013 00:33:57  startup.mdb
 1      -rw-       3792  Aug 23 2013 00:33:56  startup.cfg
 2      -rw-    53555200  Aug 23 2013 16:04:08  system.bin
 3      drw-        -   Aug 23 2013 00:03:07  seclog
```

```

4      drw-          - Aug 23 2013 00:03:07      diagfile
5      drw-          - Aug 23 2013 00:03:07      logfile
6      -rw-      9959424 Aug 23 2013 16:04:08      boot.bin
7      -rw-      9012224 Aug 21 2013 09:54:27      backup.bin

```

```
1048576 KB total (977704 KB free)
```

5. Identify the free Flash space of each subordinate switch, for example, switch 2.

```
<Sysname> dir slot2#flash:/
```

```
Directory of slot2#flash:/
```

```

0      -rw-      41424 Aug 23 2013 00:33:57      startup.mdb
1      -rw-      3792 Aug 23 2013 00:33:56      startup.cfg
2      -rw-     93871104 Aug 23 2013 16:00:08      system.bin
3      drw-          - Aug 23 2013 00:03:07      seclog
4      drw-          - Aug 23 2013 00:03:07      diagfile
5      drw-          - Aug 23 2013 00:03:07      logfile
6      -rw-     13611008 Aug 23 2013 15:59:00      boot.bin
7      -rw-      9012224 Aug 21 2013 09:54:27      backup.bin

```

```
1048576 KB total (934767 KB free)
```

6. Compare the free Flash space of each member switch with the size of the software file to load. If the space is sufficient, start the upgrade process. If not, go to the next step.
7. Delete obsolete files in Flash to free space:

CAUTION:

- To avoid data loss, do not delete the current configuration file. For information about the current configuration file, perform the **display startup** command. Hewlett Packard Enterprise recommends that you preferentially delete obsolete software images. To avoid inadvertent delete of the current software images, perform the **display boot-loader** command in any view to identify them.
- The **delete /unreserved file-url** command deletes a file permanently and the action cannot be undone.
- The **delete file-url** command moves a file to the recycle bin and the file still occupies storage space. To permanently delete the file from the recycle bin, first perform the **undelete** command to restore the file and then perform the **delete /unreserved file-url** command.

8. Delete obsolete files from the Flash memory of the master switch.

```
<Sysname> delete /unreserved flash:/backup.bin
```

```
The file cannot be restored. Delete flash:/backup.bin?[Y/N]:y
```

```
Deleting the file permanently will take a long time. Please wait...
```

```
Deleting file flash:/backup.bin...Done.
```

9. Delete obsolete files from the Flash memory of the subordinate switch.

```
<Sysname> delete /unreserved slot2#flash:/backup.bin
```

```
The file cannot be restored. Delete slot2#flash:/backup.bin?[Y/N]:y
```

```
Deleting the file permanently will take a long time. Please wait...
```

```
Deleting file slot2#flash:/backup.bin...Done.
```

Downloading software to the master switch

Before you start upgrading software images or patch packages, make sure you have downloaded the upgrading software files to the root directory in Flash memory. This section describes downloading an .ipe software file as an example.

The following are ways to download, upload, or copy files to the master switch:

- FTP download from a server
- FTP upload from a client
- TFTP download from a server
- Copying files from a USB flash drive

Prerequisites

If FTP or TFTP is used, the IRF fabric and the PC working as the FTP/TFTP server or FTP client can reach each other.

Prepare the FTP server or TFTP server program yourself for the PC. The switch series does not come with these software programs.

FTP download from a server

You can use the switch as an FTP client to download files from an FTP server.

To download a file from an FTP server, for example, the server at 10.10.110.1:

1. Run an FTP server program on the server, configure an FTP username and password, specify the working directory and copy the file, for example, **newest.ipe**, to the directory.
2. Perform the **ftp** command in user view on the IRF fabric to access the FTP server.

```
<Sysname> ftp 10.10.110.1
Trying 10.10.110.1...
Press CTRL+K to abort
Connected to 10.10.110.1
220 FTP service ready.
User(10.10.110.1:(none)):username
331 Password required for username.
Password:
230 User logged in
```

3. Enable the binary transfer mode.

```
ftp> binary
200 Type set to I.
```

4. Perform the **get** command in FTP client view to download the file from the FTP server.

```
ftp> get newest.ipe
227 Entering Passive Mode (10,10,110,1,17,97).
125 BINARY mode data connection already open, transfer starting for /newest.ipe
226 Transfer complete.
63521792 bytes received in 35 seconds (896.0 kbyte/s)
ftp> bye
221 Server closing.
```

FTP upload from a client

You can use the IRF fabric as an FTP server and upload files from a client to the IRF fabric.

To FTP upload a file from a client:

1. On the IRF fabric:
2. Enable FTP server.


```
<Sysname> system-view
[Sysname] ftp server enable
```
3. Add a local FTP user account, set its password and access service type, and assign it to the user role network-admin for uploading file to the working directory of the server.

```
[Sysname] local-user abc
```

```
[Sysname-luser-manage-abc] password simple pwd
[Sysname-luser-manage-abc] service-type ftp
[Sysname-luser-manage-abc] authorization-attribute user-role network-admin
[Sysname-luser-manage-abc] quit
[Sysname] quit
```

4. On the PC:

5. FTP to the IRF fabric (the FTP server).

```
c:\> ftp 1.1.1.1
Connected to 1.1.1.1.
220 FTP service ready.
User(1.1.1.1:(none)):abc
331 Password required for abc.
Password:
230 User logged in.
```

6. Enable the binary file transfer mode.

```
ftp> binary
200 TYPE is now 8-bit binary.
```

7. Upload the file (for example, **newest.ipe**) to the root directory in the Flash memory of the master switch.

```
ftp> put newest.ipe
200 PORT command successful
150 Connecting to port 10002
226 File successfully transferred
ftp: 63521792 bytes sent in 64.58 secs (497.60 Kbytes/sec).
```

TFTP download from a server

To download a file from a TFTP server, for example, the server at 10.10.110.1:

1. Run a TFTP server program on the server, specify the working directory, and copy the file, for example, **newest.ipe**, to the directory.
2. On the IRF fabric, perform the **tftp** command in user view to download the file to the root directory in the Flash memory of the master switch.

```
<Sysname> tftp 10.10.110.1 get newest.ipe
```

Press CTRL+C to abort.

| % Total | % Received | % Xferd | Average Speed | Time | Time | Time | Current |
|-----------|------------|---------|---------------|----------|---------|----------|---------|
| | | | Dload Upload | Total | Spent | Left | Speed |
| 100 60.5M | 0 60.5M | 0 0 | 143k 0 | --:--:-- | 0:03:38 | --:--:-- | 142k |

Copying files from a USB flash drive

Every 5940 switch provides a USB port for you to copy files from a USB flash drive.

To copy a file from a USB flash drive to the Flash memory of the master switch:

1. Plug the USB flash drive in the USB port of the switch.
2. Copy the file (for example, **newest.ipe**) to the Flash memory of the switch.

```
<Sysname> cd usba:
```

```
<Sysname> copy usba:/newest.ipe newest.ipe
```

```
Copy usba:/newest.ipe to flash:/newest.ipe?[Y/N]:y
```

```
Start to copy usba:/newest.ipe to flash:/newest.ipe... Done.
```

Upgrading the software images

To upgrade the software images:

1. Specify the upgrading image file (**newest.ipe** in this example) used at the next startup for the master switch, and assign the M attribute to the boot and system images in the file.

```
<Sysname> boot-loader file flash:/newest.ipe slot 1 main
```

```
Verifying image file....Done.
```

```
Images in IPE:
```

```
boot.bin
```

```
system.bin
```

```
This command will set the main startup software images. Continue? [Y/N]:y
```

```
Add images to target slot.
```

```
The images that have passed all examinations will be used as the main startup software images at the next reboot on slot 1.
```

2. Specify the upgrading image file used at next startup for the subordinate switch, and assign the M attribute to the boot and system images in the file. (As a result, the subordinate switch automatically copies the file to the root directory in its Flash memory.)

```
<Sysname> boot-loader file flash:/newest.ipe slot 2 main
```

```
Verifying image file....Done.
```

```
Images in IPE:
```

```
boot.bin
```

```
system.bin
```

```
This command will set the main startup software images. Continue? [Y/N]:y
```

```
Add images to target slot.
```

```
The images that have passed all examinations will be used as the main startup software images at the next reboot on slot 2.
```

3. (Optional) If the IRF fabric size has a lot of members, enable the software auto-update function.

```
<Sysname> system-view
```

```
[Sysname] irf auto-update enable
```

```
[Sysname] quit
```

Software auto-update is typically used for synchronizing the software images of the master switch to new member switches when you expand the IRF fabric. This function enables a subordinate switch to compare its main startup software image version with that of the IRF master. If the versions are different, the subordinate switch automatically downloads the current software images from the master, sets the downloaded images as the main software images at the next reboot, and automatically reboots with the new images to re-join the IRF fabric. In this upgrade process, the function avoids the failure of assign all the subordinate switch the same main software image file as the master switch causing an upgrade failure.

4. Save the current configuration in any view to prevent data loss.

```
<Sysname> save
```

```
The current configuration will be written to the device. Are you sure? [Y/N]:y
```

```
Please input the file name(*.cfg)[flash:/startup.cfg]
```

```
(To leave the existing filename unchanged, press the enter key):
```

```
flash:/startup.cfg exists, overwrite? [Y/N]:y
```

```
Validating file. Please wait.....
```

```
Saved the current configuration to mainboard device successfully.
```

```
Slot 2:
```

```
Save next configuration file successfully.
```

5. Reboot the IRF fabric to complete the upgrade.

```
<Sysname> reboot
```

```
Start to check configuration with next startup configuration file, please wait.
```

```
.....DONE!
```

```
This command will reboot the device. Continue? [Y/N]:y
```

```
Now rebooting, please wait...
```

The system automatically loads the .bin boot and system images in the .ipe file and sets them as the startup software images.

6. Perform the **display version** command in any view to verify that the current main software images have been updated (details not shown).

NOTE:

The system automatically checks the compatibility of the Boot ROM image and the boot and system images during the reboot. If you are prompted that the Boot ROM image in the upgrading image file is different than the current Boot ROM image, upgrade both the basic and extended sections of the Boot ROM image for compatibility. If you choose to not upgrade the Boot ROM image, the system will ask for an upgrade at the next reboot performed by powering on the switch or rebooting from the CLI (promptly or as scheduled). If you fail to make any choice in the required time, the system upgrades the entire Boot ROM image.

Installing a patch package

To install a patch package, for example, **system-patch.bin**:

1. Activate the patch package on the master switch and the subordinate switch.

```
<Sysname> install activate patch flash:/system-patch.bin slot 1
```

```
<Sysname> install activate patch flash:/system-patch.bin slot 2
```

2. Verify that the patch package has been activated.

```
<Sysname> display install active
```

```
Active packages on slot 1:
```

```
flash:/boot.bin
```

```
flash:/system.bin
```

```
flash:/system-patch.bin
```

```
Active packages on slot 2:
```

```
flash:/boot.bin
```

```
flash:/system.bin
```

```
flash:/system-patch.bin
```

3. Commit the installation so the patch package continues to take effect after a reboot.

```
<Sysname> install commit
```

4. Verify that the patch package installation has been committed.

```
<Sysname> display install committed
```

```
Committed packages on slot 1:
```

```
flash:/boot.bin
```

```
flash:/system.bin
```

```
flash:/system-patch.bin
```

```
Committed packages on slot 2:
```

```
flash:/boot.bin
```

```
flash:/system.bin
```

```
flash:/system-patch.bin
```

For more information about installing patch packages, see HPE FlexFabric 5940 Switch Series Fundamentals Configuration Guide.

Upgrading from the Boot menu

You can upgrade the Boot ROM image and software images but not patch packages from the Boot menu.

In this approach, you must access the Boot menu of each member switch to upgrade their software one by one. If you are upgrading software images for an IRF fabric, using the CLI is a better choice.

The following sections describe the methods of upgrading software images:

- [Using TFTP to upgrade software images through the management Ethernet port](#)
- [Using FTP to upgrade software through the management Ethernet port](#)
- [Using XMODEM to upgrade software through the console port](#)

The following sections describe the methods of upgrading Boot ROM images:

- [Using TFTP to upgrade Boot ROM through the management Ethernet port](#)
- [Using FTP to upgrade Boot ROM through the management Ethernet port](#)
- [Using XMODEM to upgrade Boot ROM through the console port](#)



TIP:

Upgrading through an Ethernet port is faster than through the console port.

Prerequisites

Make sure that the prerequisites are met before you start upgrading software from the Boot menu.

Upgrading environment

Use a console cable to connect the console terminal, for example, a PC, to the console port on the switch. Run a terminal emulator program on the console terminal and set the following terminal settings:

- **Bits per second**—9,600
- **Data bits**—8
- **Parity**—None
- **Stop bits**—1
- **Flow control**—None
- **Emulation**—VT100

TFTP/FTP download

To use TFTP or FTP:

- Run a TFTP or FTP server program on the file server or the console terminal.
- Copy the upgrade file to the file server.
- Correctly set the working directory on the TFTP or FTP server.
- Make sure that the file server and the switch can reach each other.

Storage space

Make sure that sufficient space is available for the upgrading software file. If no sufficient space is available, delete obsolete files as described in "[Managing files from the Boot menu](#)."

Upgrading time

Make sure that the upgrade has minimal impact on the network services. During the upgrade, the switch cannot provide any services.

Accessing the Boot menu

Power on the switch (for example, an HPE FF 5940-32QSFP+ Switch), and you can see the following information:

Starting.....

Press Ctrl+D to access BASIC BOOT MENU...

Press Ctrl+T to start heavy memory test

```
*****
*
*          HPE FF 5940-32QSFP+ Switch BOOTROM, Version 205          *
*
*****
Copyright (c) 2010-2016 Hewlett-Packard Development Company, L.P.
```

Creation Date : Jan 6 2013, 14:25:58

CPU Clock Speed : 1000MHz

Memory Size : 4096MB

Flash Size : 1024MB

CPLD Version : 002/002

PCB Version : Ver.A

Mac Address : 00E0FC005800

Press Ctrl+B to access EXTENDED BOOT MENU...1

Press one of the shortcut key combinations at prompt.

Table 8 Shortcut keys

| Shortcut keys | Prompt message | Function | Remarks |
|---------------|---|----------------------------------|---|
| Ctrl+B | Press Ctrl+B to enter Extended Boot menu... | Accesses the extended Boot menu. | Press the keys within 1 second (in fast startup mode) or 5 seconds (in full startup mode) after the message appears. You can upgrade and manage system software and Boot ROM from this menu. |
| Ctrl+D | Press Ctrl+D to access BASIC BOOT MENU | Accesses the basic Boot menu. | Press the keys within 1 seconds after the message appears. You can upgrade Boot ROM or access the extended Boot ROM segment from this menu. |
| Ctrl+T | Press Ctrl+T to start heavy memory test | Performs a RAM pressure test. | Press the keys within 1 second after the message appears. |

Accessing the basic Boot menu

If the extended Boot ROM segment has corrupted, you can repair or upgrade it from the basic Boot menu.

Press **Ctrl+D** within 1 seconds after the "Press Ctrl+D to access BASIC BOOT MENU" prompt message appears. If you fail to do this within the time limit, the system starts to run the extended Boot ROM segment.

```
*****
*
*                                     BASIC BOOTROM, Version 205
*
*****
```

BASIC BOOT MENU

```
1. Update full BootRom
2. Update extended BootRom
3. Update basic BootRom
4. Boot extended BootRom
0. Reboot
Ctrl+U: Access BASIC ASSISTANT MENU
```

Enter your choice(0-4):

Table 9 Basic Boot ROM menu options

| Option | Task |
|-------------------------------------|---|
| 1. Update full BootRom | Update the entire Boot ROM, including the basic segment and the extended segment. To do so, you must use XMODEM and the console port. For more information, see Using XMODEM to upgrade Boot ROM through the console port . |
| 2. Update extended BootRom | Update the extended Boot ROM segment. To do so, you must use XMODEM and the console port. For more information, see Using XMODEM to upgrade Boot ROM through the console port . |
| 3. Update basic BootRom | Update the basic Boot ROM segment. To do so, you must use XMODEM and the console port. For more information, see Using XMODEM to upgrade Boot ROM through the console port . |
| 4. Boot extended BootRom | Access the extended Boot ROM segment. For more information, see Accessing the extended Boot menu . |
| 0. Reboot | Reboot the switch. |
| Ctrl+U: Access BASIC ASSISTANT MENU | Press Ctrl + U to access the BASIC ASSISTANT menu (see Table 10). |

Table 10 BASIC ASSISTANT menu options

| Option | Task |
|-------------|--------------------------|
| 1. RAM Test | Perform a RAM self-test. |

| Option | Task |
|------------------------|--------------------------------|
| 0. Return to boot menu | Return to the basic Boot menu. |

Accessing the extended Boot menu

Press **Ctrl+B** within 1 second (in fast startup mode) or 5 seconds (in full startup mode) after the "Press Ctrl-B to enter Extended Boot menu..." prompt message appears. If you fail to do this, the system starts decompressing the system software.

Alternatively, you can enter 4 in the basic Boot menu to access the extended Boot menu.

The "Password recovery capability is enabled." or "Password recovery capability is disabled." message appears, followed by the extended Boot menu. Availability of some menu options depends on the state of password recovery capability (see [Table 11](#)). For more information about password recovery capability, see *HPE FlexFabric 5940 Switch Series Fundamentals Configuration Guide*.

Password recovery capability is enabled.

```

EXTENDED BOOT MENU

1. Download image to flash
2. Select image to boot
3. Display all files in flash
4. Delete file from flash
5. Restore to factory default configuration
6. Enter BootRom upgrade menu
7. Skip current system configuration
8. Set switch startup mode
0. Reboot
Ctrl+Z: Access EXTENDED ASSISTANT MENU
Ctrl+F: Format file system
Ctrl+P: Change authentication for console login
Ctrl+Y: Change Work Mode
Ctrl+R: Download image to SDRAM and run

Enter your choice(0-8):

```

Table 11 Extended Boot ROM menu options

| Option | Tasks |
|-------------------------------|--|
| 1. Download image to flash | Download a software image file to the flash. |
| 2. Select image to boot | <ul style="list-style-type: none"> Specify the main and backup software image file for the next startup. Specify the main and backup configuration files for the next startup. This task can be performed only if password recovery capability is enabled. |
| 3. Display all files in flash | Display files on the flash. |
| 4. Delete file from flash | Delete files to free storage space. |

| Option | Tasks |
|---|---|
| 5. Restore to factory default configuration | Delete the current next-startup configuration files and restore the factory-default configuration. This option is available only if password recovery capability is disabled. |
| 6. Enter BootRom upgrade menu | Access the Boot ROM upgrade menu. |
| 7. Skip current system configuration | Start the switch without loading any configuration file. This is a one-time operation and takes effect only for the first system boot or reboot after you choose this option. This option is available only if password recovery capability is enabled. |
| 8. Set switch startup mode | Set the startup mode to fast startup mode or full startup mode. |
| 0. Reboot | Reboot the switch. |
| Ctrl+F: Format file system | Format the current storage medium. |
| Ctrl+P: Change authentication for console login | Skip the authentication for console login. This is a one-time operation and takes effect only for the first system boot or reboot after you choose this option. This option is available only if password recovery capability is enabled. |
| Ctrl+Y: Change the operating mode | If PEX mode is enabled, the operation disables PEX mode. If PEX mode is disabled, the operation enables PEX mode. |
| Ctrl+R: Download image to SDRAM and run | Download a system software image and start the switch with the image. This option is available only if password recovery capability is enabled. |
| Ctrl+Z: Access EXTENDED ASSISTANT MENU | Access the EXTENDED ASSISTANT MENU. For options in the menu, see Table 12 . |

Table 12 EXTENDED ASSISTANT menu options

| Option | Task |
|------------------------|--|
| 1. Display Memory | Display data in the memory. |
| 2. Search Memory | Search the memory for a specific data segment. |
| 0. Return to boot menu | Return to the extended Boot ROM menu. |

Using TFTP to upgrade software images through the management Ethernet port

1. Enter **1** in the Boot menu to access the file transfer protocol submenu.

1. Set TFTP protocol parameters
2. Set FTP protocol parameters
3. Set XMODEM protocol parameters
0. Return to boot menu

Enter your choice(0-3):

2. Enter **1** to set the TFTP parameters.

```
Load File Name      :update.ipe
Server IP Address   :192.168.0.3
Local IP Address    :192.168.0.2
```

Subnet Mask :255.255.255.0
Gateway IP Address :0.0.0.0

Table 13 TFTP parameter description

| Item | Description |
|--------------------|---|
| Load File Name | Name of the file to download (for example, update.ipe). |
| Server IP Address | IP address of the TFTP server (for example, 192.168.0.3). |
| Local IP Address | IP address of the switch (for example, 192.168.0.2). |
| Subnet Mask | Subnet mask of the switch (for example, 255.255.255.0). |
| Gateway IP Address | IP address of the gateway (in this example, no gateway is required because the server and the switch are on the same subnet). |

NOTE:

- To use the default setting for a field, press **Enter** without entering any value.
- If the switch and the server are on different subnets, you must specify a gateway address for the switch.

3. Enter all required parameters, and enter **Y** to confirm the settings. The following prompt appears:

Are you sure to download file to flash? Yes or No (Y/N):Y

4. Enter **Y** to start downloading the image file. To return to the Boot menu, enter **N**.

```
Loading.....
.....
.....Done!
```

5. Enter the **M** (main), **B** (backup), or **N** (none) attribute for the images. In this example, assign the main attribute to the images.

Please input the file attribute (Main/Backup/None) M

Image file boot.bin is self-decompressing...

Free space: 534980608 bytes

```
Writing flash.....
.....Done!
```

Image file system.bin is self-decompressing...

Free space: 525981696 bytes

```
Writing flash.....
.....
.....
.....
.....Done!
```

NOTE:

- The switch always attempts to boot first with the main images, and if the attempt fails, for example, because the main images are not available, the switch tries to boot with the backup images. An image with the none attribute is just stored in Flash memory for backup and you must change its attribute to make it usable at reboot.
 - If an image with the same attribute as the image you are loading is already in Flash memory, the attribute of the old image changes to none after the new image becomes valid.
-

6. Enter 0 in the Boot menu to reboot the switch with the new software images.

EXTENDED BOOT MENU

```
1. Download image to flash
2. Select image to boot
3. Display all files in flash
4. Delete file from flash
5. Restore to factory default configuration
6. Enter BootRom upgrade menu
7. Skip current system configuration
8. Set switch startup mode
0. Reboot
Ctrl+Z: Access EXTENDED ASSISTANT MENU
Ctrl+F: Format file system
Ctrl+P: Change authentication for console login
Ctrl+Y: Change Work Mode
Ctrl+R: Download image to SDRAM and run
```

Enter your choice(0-8):0

Using FTP to upgrade software through the management Ethernet port

1. Enter 1 in the Boot menu to access the file transfer protocol submenu.

```
1. Set TFTP protocol parameters
2. Set FTP protocol parameters
3. Set XMODEM protocol parameters
0. Return to boot menu
```

Enter your choice(0-3):

2. Enter 2 to set the FTP parameters.

```
Load File Name      :update.ipe
Server IP Address   :192.168.0.3
Local IP Address    :192.168.0.2
Subnet Mask         :255.255.255.0
Gateway IP Address  :0.0.0.0
FTP User Name       :switch
FTP User Password   :***
```

Table 14 FTP parameter description

| Item | Description |
|--------------------|---|
| Load File Name | Name of the file to download (for example, update.ipe). |
| Server IP Address | IP address of the FTP server (for example, 192.168.0.3). |
| Local IP Address | IP address of the switch (for example, 192.168.0.2). |
| Subnet Mask | Subnet mask of the switch (for example, 255.255.255.0). |
| Gateway IP Address | IP address of the gateway (in this example, no gateway is required because the server and the switch are on the same subnet). |
| FTP User Name | Username for accessing the FTP server, which must be the same as configured on the FTP server. |
| FTP User Password | Password for accessing the FTP server, which must be the same as configured on the FTP server. |

NOTE:

- To use the default setting for a field, press **Enter** without entering any value.
- If the switch and the server are on different subnets, you must specify a gateway address for the switch.

3. Enter all required parameters, and enter **Y** to confirm the settings. The following prompt appears:

```
Are you sure to download file to flash? Yes or No (Y/N):Y
```

4. Enter **Y** to start downloading the image file. To return to the Boot menu, enter **N**.

```
Loading.....
.....
.....
.....Done!
```

5. Enter the **M** (main), **B** (backup), or **N** (none) attribute for the images. In this example, assign the main attribute to the images.

```
Please input the file attribute (Main/Backup/None) M
Image file boot.bin is self-decompressing...
Free space: 534980608 bytes
Writing flash.....
.....
.....Done!
Image file system.bin is self-decompressing...
Free space: 525981696 bytes
Writing flash.....
.....
.....
.....
.....
.....
.....Done!
```

EXTENDED BOOT MENU

1. Download image to flash
2. Select image to boot

```

3. Display all files in flash
4. Delete file from flash
5. Restore to factory default configuration
6. Enter BootRom upgrade menu
7. Skip current system configuration
8. Set switch startup mode
0. Reboot

Ctrl+Z: Access EXTENDED ASSISTANT MENU
Ctrl+F: Format file system
Ctrl+P: Change authentication for console login
Ctrl+Y: Change Work Mode
Ctrl+R: Download image to SDRAM and run

```

```
Enter your choice(0-8):0
```

NOTE:

- The switch always attempts to boot first with the main images, and if the attempt fails, for example, because the main images not available, the switch tries to boot with the backup images. An image with the none attribute is just stored in Flash memory for backup and you must change its attribute to make it usable at reboot.
 - If an image with the same attribute as the image you are loading is already in Flash memory, the attribute of the old image changes to none after the new image becomes valid.
-

6. Enter **0** in the Boot menu to reboot the switch with the new software images.

Using XMODEM to upgrade software through the console port

XMODEM download through the console port is slower than TFTP or FTP download through the management Ethernet port. To save time, use the management Ethernet port as long as possible.

1. Enter **1** in the Boot menu to access the file transfer protocol submenu.

```

1. Set TFTP protocol parameters
2. Set FTP protocol parameters
3. Set XMODEM protocol parameters
0. Return to boot menu

```

```
Enter your choice(0-3):
```

2. Enter **3** to set the XMODEM download baud rate.

```
Please select your download baudrate:
```

```

1.* 9600
2. 19200
3. 38400
4. 57600
5. 115200
0. Return to boot menu

```

```
Enter your choice(0-5):5
```

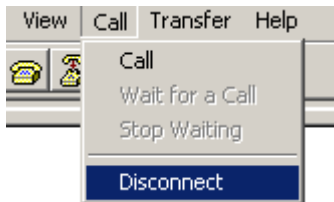
3. Select an appropriate download rate, for example, enter **5** to select 115200 bps.

```
Download baudrate is 115200 bps
```

Please change the terminal's baudrate to 115200 bps and select XMODEM protocol
Press enter key when ready

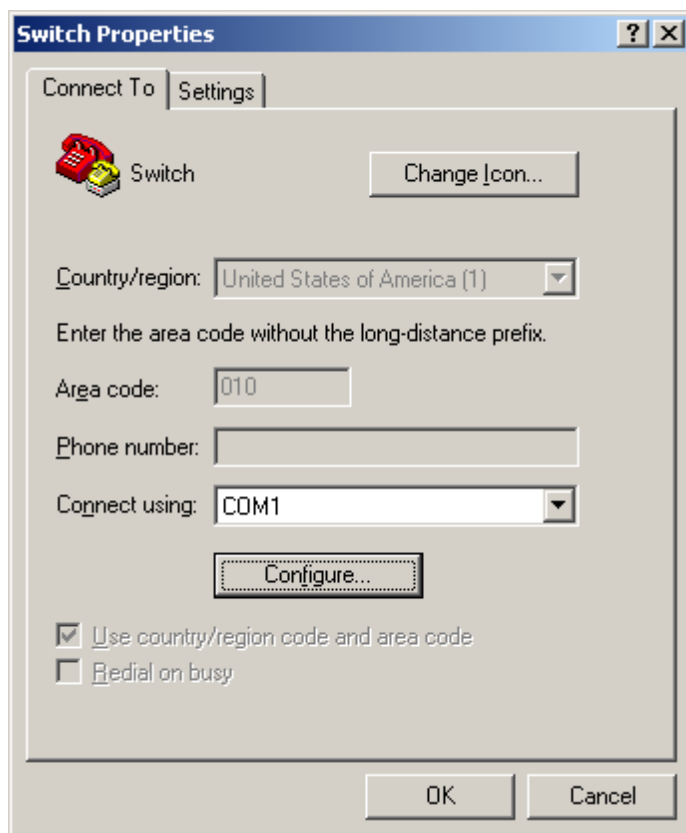
4. Set the serial port on the terminal to use the same baud rate and protocol as the console port. If you select 9600 bps as the download rate for the console port, skip this task.
5. Select **Call > Disconnect** in the HyperTerminal window to disconnect the terminal from the switch.

Figure 2 Disconnecting the terminal from the switch



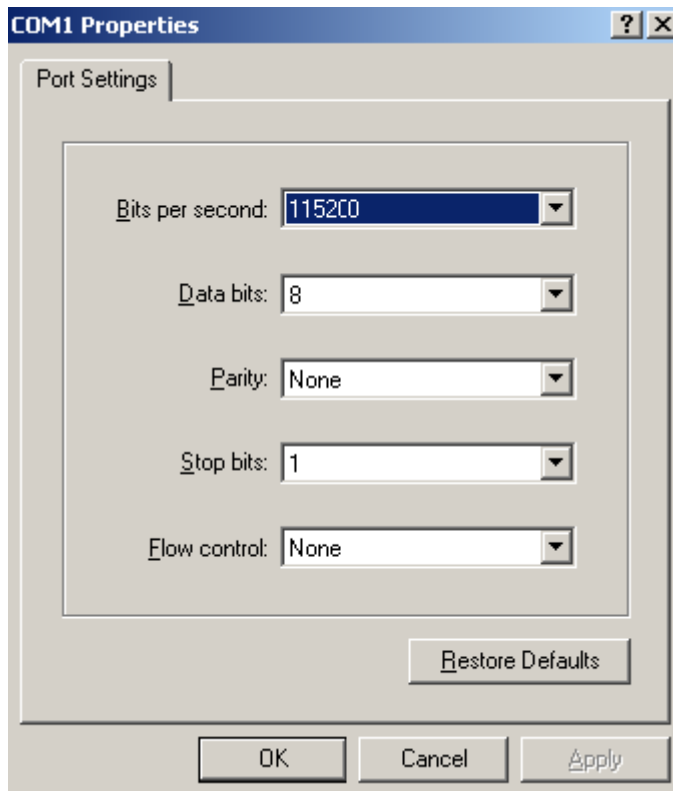
6. Select **File > Properties**, and in the **Properties** dialog box, click **Configure**.

Figure 3 Properties dialog box



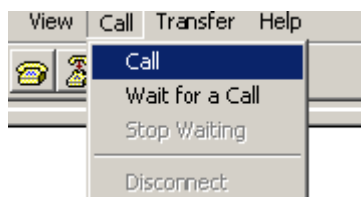
7. Select **115200** from the **Bits per second** list and click **OK**.

Figure 4 Modifying the baud rate



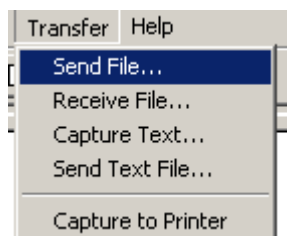
8. Select **Call** > **Call** to reestablish the connection.

Figure 5 Reestablishing the connection



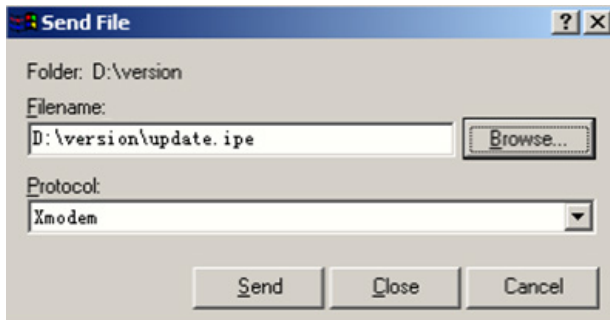
9. Press **Enter**. The following prompt appears:
Are you sure to download file to flash? Yes or No (Y/N):Y
10. Enter **Y** to start downloading the file. (To return to the Boot menu, enter **N**.)
Now please start transfer file with XMODEM protocol
If you want to exit, Press <Ctrl+X>
Loading ...CCCCCCCCCCCCCCCCCCCCCCCCCCCC
11. Select **Transfer** > **Send File** in the HyperTerminal window.

Figure 6 Transfer menu



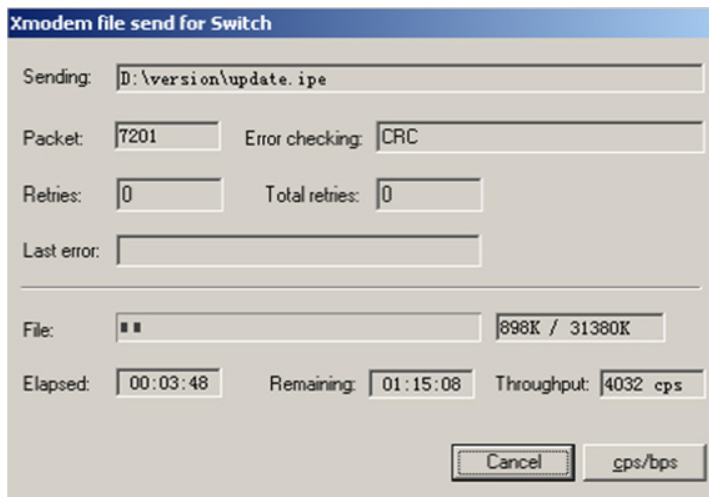
12. In the dialog box that appears, click **Browse** to select the source file, and select **Xmodem** from the **Protocol** list.

Figure 7 File transmission dialog box



13. Click **Send**. The following dialog box appears:

Figure 8 File transfer progress



14. Enter the **M** (main), **B** (backup), or **N** (none) attribute for the images. In this example, assign the main attribute to the images.

Please input the file attribute (Main/Backup/None) m

The boot.bin image is self-decompressing...

At the **Load File name** prompt, enter a name for the Boot image to be saved to Flash memory.

Load File name : default_file boot-update.bin

Free space: 470519808 bytes

Writing flash.....
.....Done!

The system-update.bin image is self-decompressing...

At the **Load File name** prompt, enter a name for the system image to be saved to Flash memory.

Load File name : default_file system-update.bin

Free space: 461522944 bytes

Writing flash.....
.....Done!

Your baudrate should be set to 9600 bps again!

Press enter key when ready

NOTE:

- The switch always attempts to boot first with the main images, and if the attempt fails, for example, because the main images are not available, the switch tries to boot with the backup images. An image with the none attribute is just stored in Flash memory for backup and you must change its attribute to make it usable at reboot.
 - If an image with the same attribute as the image you are loading is already in Flash memory, the attribute of the old image changes to none after the new image becomes valid.
-

- 15.** If the baud rate of the HyperTerminal is not 9600 bps, restore it to 9600 bps. If the baud rate is 9600 bps, skip this step.

To access the switch through the console port after a reboot, you must perform this step, because the console port rate reverts to 9600 bps at a reboot.

- 16.** Press **Enter** to access the Boot menu.

EXTENDED BOOT MENU

```
1. Download image to flash
2. Select image to boot
3. Display all files in flash
4. Delete file from flash
5. Restore to factory default configuration
6. Enter BootRom upgrade menu
7. Skip current system configuration
8. Set switch startup mode
0. Reboot

Ctrl+Z: Access EXTENDED ASSISTANT MENU
Ctrl+F: Format file system
Ctrl+P: Change authentication for console login
Ctrl+Y: Change Work Mode
Ctrl+R: Download image to SDRAM and run
```

```
Enter your choice(0-8):0
```

- 17.** Enter **0** to reboot the system with the new software images.

Using TFTP to upgrade Boot ROM through the management Ethernet port

- 1.** Enter **6** in the Boot menu to access the Boot ROM update menu.

```
1. Update full BootRom
2. Update extended BootRom
3. Update basic BootRom
0. Return to boot menu
```

```
Enter your choice(0-3):
```

- 2.** Enter **1** in the Boot ROM update menu to upgrade the full Boot ROM.

The file transfer protocol submenu appears:

```
1. Set TFTP protocol parameters
2. Set FTP protocol parameters
3. Set XMODEM protocol parameters
```

0. Return to boot menu

Enter your choice(0-3):

3. Enter 1 to set the TFTP parameters.

Load File Name :update.btm
Server IP Address :192.168.0.3
Local IP Address :192.168.0.2
Subnet Mask :255.255.255.0
Gateway IP Address :0.0.0.0

Table 15 TFTP parameter description

| Item | Description |
|--------------------|---|
| Load File Name | Name of the file to download (for example, update.btm). |
| Server IP Address | IP address of the TFTP server (for example, 192.168.0.3). |
| Local IP Address | IP address of the switch (for example, 192.168.0.2). |
| Subnet Mask | Subnet mask of the switch (for example, 255.255.255.0). |
| Gateway IP Address | IP address of the gateway (in this example, no gateway is required because the server and the switch are on the same subnet). |

NOTE:

- To use the default setting for a field, press **Enter** without entering any value.
- If the switch and the server are on different subnets, you must specify a gateway address for the switch.
- If .bin files are used for upgrade, specify the .bin files in the order of the boot image, system image, and feature images. If you specify a .bin file for a feature image before the .bin file for the system image or the boot image, the upgrade might fail

4. Enter all required parameters and press **Enter to start downloading the file.**

Loading.....Done!

5. Enter Y at the prompt to upgrade the basic Boot ROM section.

Will you Update Basic BootRom? (Y/N):Y
Updating Basic BootRom.....Done.

6. Enter Y at the prompt to upgrade the extended Boot ROM section.

Updating extended BootRom? (Y/N):Y
Updating extended BootRom.....Done.

7. Enter 0 in the Boot ROM update menu to return to the Boot menu.

1. Update full BootRom
2. Update extended BootRom
3. Update basic BootRom
0. Return to boot menu

Enter your choice(0-3):

8. Enter 0 in the Boot menu to reboot the switch with the new Boot ROM image.

Using FTP to upgrade Boot ROM through the management Ethernet port

1. Enter **6** in the Boot menu to access the Boot ROM update menu.

1. Update full BootRom
2. Update extended BootRom
3. Update basic BootRom
0. Return to boot menu

Enter your choice(0-3):

2. Enter 1 in the Boot ROM update menu to upgrade the full Boot ROM.

The file transfer protocol submenu appears:

1. Set TFTP protocol parameters
2. Set FTP protocol parameters
3. Set XMODEM protocol parameters
0. Return to boot menu

Enter your choice(0-3):

3. Enter 2 to set the FTP parameters.

```
Load File Name      :update.btm
Server IP Address   :192.168.0.3
Local IP Address    :192.168.0.2
Subnet Mask         :255.255.255.0
Gateway IP Address  :0.0.0.0
FTP User Name       :switch
FTP User Password   :***
```

Table 16 FTP parameter description

| Item | Description |
|--------------------|---|
| Load File Name | Name of the file to download (for example, update.btm). |
| Server IP Address | IP address of the FTP server (for example, 192.168.0.3). |
| Local IP Address | IP address of the switch (for example, 192.168.0.2). |
| Subnet Mask | Subnet mask of the switch (for example, 255.255.255.0). |
| Gateway IP Address | IP address of the gateway (in this example, no gateway is required because the server and the switch are on the same subnet). |
| FTP User Name | Username for accessing the FTP server, which must be the same as configured on the FTP server. |
| FTP User Password | Password for accessing the FTP server, which must be the same as configured on the FTP server. |

NOTE:

- To use the default setting for a field, press **Enter** without entering any value.
- If the switch and the server are on different subnets, you must specify a gateway address for the switch.
- If .bin files are used for upgrade, specify the .bin files in the order of the boot image, system image, and feature images. If you specify a .bin file for a feature image before the .bin file for the system image or the boot image, the upgrade might fail

4. Enter all required parameters and press **Enter** to start downloading the file.

```
Loading.....Done!
```

5. Enter Y at the prompt to upgrade the basic Boot ROM section.

```
Will you Update Basic BootRom? (Y/N):Y
```

- ```
Updating Basic BootRom.....Done.
```
6. Enter **Y** at the prompt to upgrade the extended Boot ROM section.
 

```
Updating extended BootRom? (Y/N):Y
Updating extended BootRom.....Done.
```
  7. Enter **0** in the Boot ROM update menu to return to the Boot menu.
 

```
1. Update full BootRom
2. Update extended BootRom
3. Update basic BootRom
0. Return to boot menu

Enter your choice(0-3):
```
  8. Enter **0** in the Boot menu to reboot the switch with the new Boot ROM image.

## Using XMODEM to upgrade Boot ROM through the console port

XMODEM download through the console port is slower than TFTP or FTP download through the management Ethernet port. To save time, use the management Ethernet port as long as possible.

1. Enter **6** in the Boot menu to access the Boot ROM update menu.
 

```
1. Update full BootRom
2. Update extended BootRom
3. Update basic BootRom
0. Return to boot menu

Enter your choice(0-3):
```
2. Enter **1** in the Boot ROM update menu to upgrade the full Boot ROM.

The file transfer protocol submenu appears:

- ```
1. Set TFTP protocol parameters
2. Set FTP protocol parameters
3. Set XMODEM protocol parameters
0. Return to boot menu

Enter your choice(0-3):
```
3. Enter **3** to set the XMODEM download baud rate.

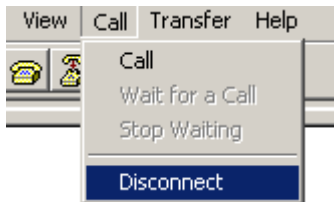

```
Please select your download baudrate:
1.* 9600
2. 19200
3. 38400
4. 57600
5. 115200
0. Return to boot menu

Enter your choice(0-5):5
```
 4. Select an appropriate download rate, for example, enter **5** to select 115200 bps.


```
Download baudrate is 115200 bps
Please change the terminal's baudrate to 115200 bps and select XMODEM protocol
Press enter key when ready
```

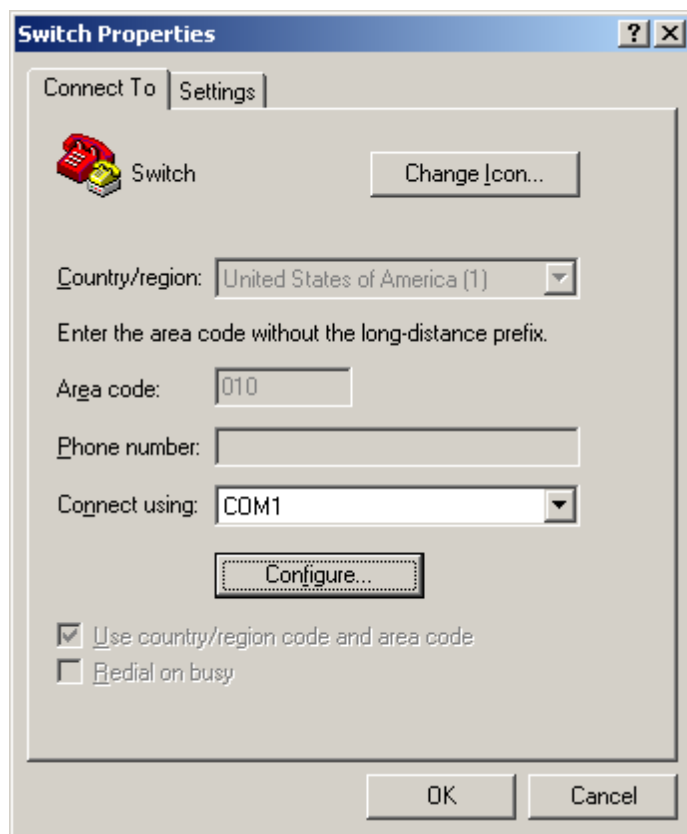
5. Set the serial port on the terminal to use the same baud rate and protocol as the console port. If you select 9600 bps as the download rate for the console port, skip this task.
6. Select **Call > Disconnect** in the HyperTerminal window to disconnect the terminal from the switch.

Figure 9 Disconnecting the terminal from the switch



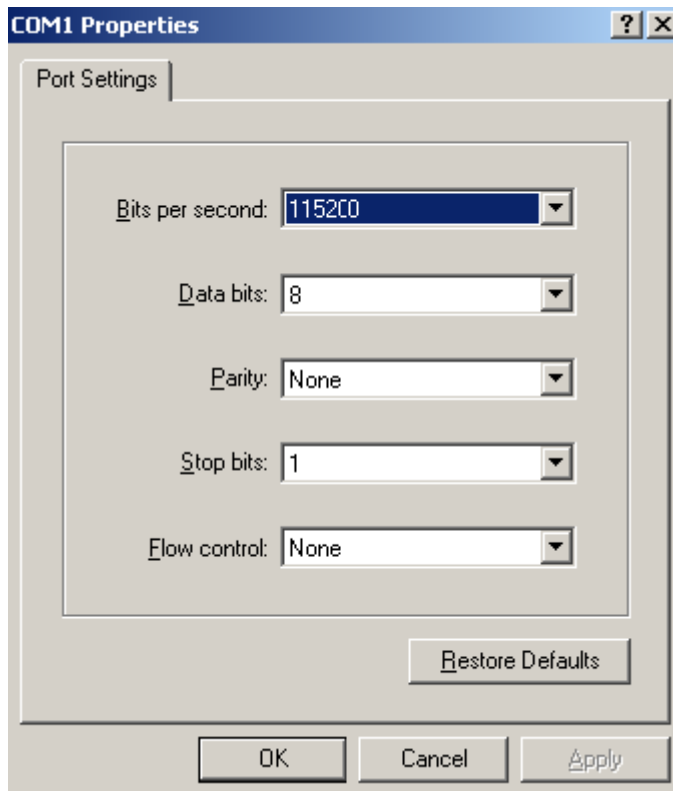
7. Select **File > Properties**, and in the **Properties** dialog box, click **Configure**.

Figure 10 Properties dialog box



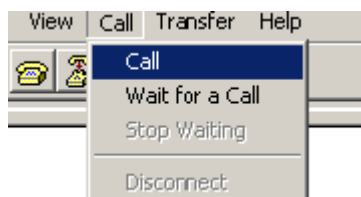
8. Select **115200** from the **Bits per second** list and click **OK**.

Figure 11 Modifying the baud rate



9. Select **Call > Call** to reestablish the connection.

Figure 12 Reestablishing the connection

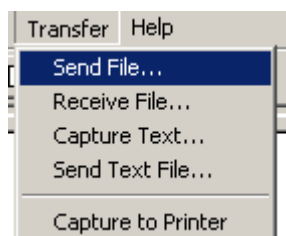


10. Press **Enter** to start downloading the file.

```
Now please start transfer file with XMODEM protocol
If you want to exit, Press <Ctrl+X>
Loading ...CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
```

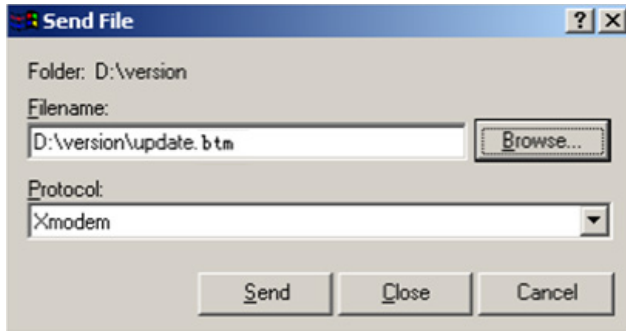
11. Select **Transfer > Send File** in the HyperTerminal window.

Figure 13 Transfer menu



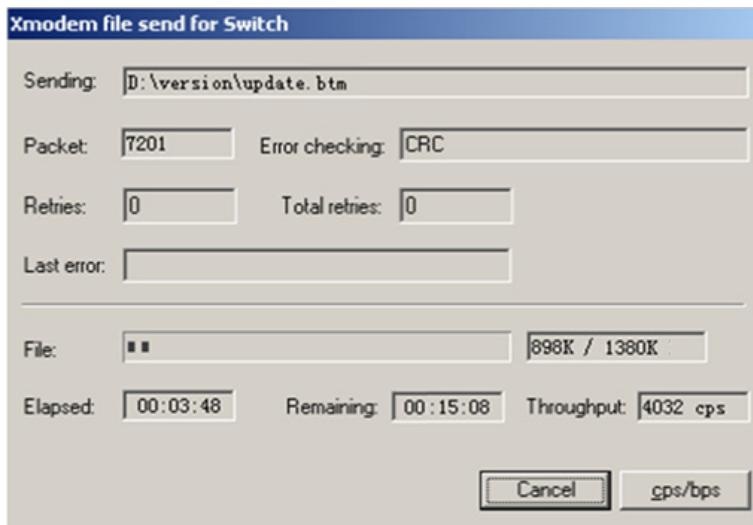
12. In the dialog box that appears, click **Browse** to select the source file, and select **Xmodem** from the **Protocol** list.

Figure 14 File transmission dialog box



13. Click **Send**. The following dialog box appears:

Figure 15 File transfer progress



14. Enter **Y** at the prompt to upgrade the basic Boot ROM section.

```
Loading ...CCCCCCCCCCCCCCCC ...Done!  
Will you Update Basic BootRom? (Y/N):Y  
Updating Basic BootRom.....Done.
```

15. Enter **Y** at the prompt to upgrade the extended Boot ROM section.

```
Updating extended BootRom? (Y/N):Y  
Updating extended BootRom.....Done.
```

16. If the baud rate of the HyperTerminal is not 9600 bps, restore it to 9600 bps at the prompt. If the baud rate is 9600 bps, skip this step.

Please change the terminal's baudrate to 9600 bps, press ENTER when ready.

To access the switch through the console port after a reboot, you must perform this step, because the console port rate reverts to 9600 bps at a reboot.

17. Press **Enter** to access the Boot ROM update menu.

18. Enter **0** in the Boot ROM update menu to return to the Boot menu.

```
1. Update full BootRom  
2. Update extended BootRom  
3. Update basic BootRom  
0. Return to boot menu
```

Enter your choice(0-3):

19. Enter **0** in the Boot menu to reboot the switch with the new Boot ROM image.

Managing files from the Boot menu

From the Boot menu, you can display files in Flash memory to check for obsolete files, incorrect files, or space insufficiency, delete files to release storage space, or change the attributes of software images.

Displaying all files

Enter **3** in the Boot menu to display all files in Flash memory and identify the free space size.

```
EXTENDED BOOT MENU

1. Download image to flash
2. Select image to boot
3. Display all files in flash
4. Delete file from flash
5. Restore to factory default configuration
6. Enter BootRom upgrade menu
7. Skip current system configuration
8. Set switch startup mode
0. Reboot

Ctrl+Z: Access EXTENDED ASSISTANT MENU
Ctrl+F: Format file system
Ctrl+P: Change authentication for console login
Ctrl+Y: Change Work Mode
Ctrl+R: Download image to SDRAM and run
```

Enter your choice(0-8): 3

The following is a sample output:

Display all file(s) in flash:

| File Number | File Size(bytes) | File Name |
|-------------|------------------|----------------------------|
| 1 | 8177 | flash:/testbackup.cfg |
| 2(*) | 53555200 | flash:/system.bin |
| 3(*) | 9959424 | flash:/boot.bin |
| 4 | 3678 | flash:/startup.cfg_backup |
| 5 | 30033 | flash:/default.mdb |
| 6 | 42424 | flash:/startup.mdb |
| 7 | 18 | flash:/pathfile |
| 8 | 232311 | flash:/logfile/logfile.log |
| 9 | 5981 | flash:/startup.cfg_back |
| 10(*) | 6098 | flash:/startup.cfg |
| 11 | 20 | flash:/snmpboots |

Free space: 1009906637 bytes
The current image is boot.bin
(*)-with main attribute

(b)-with backup attribute
(*b)-with both main and backup attribute

Deleting files

If storage space is insufficient, delete obsolete files to free up storage space.

To delete files:

1. Enter **4** in the Boot menu:

Deleting the file in flash:

| File Number | File Size(bytes) | File Name |
|-------------|------------------|----------------------------|
| 1 | 8177 | flash:/testbackup.cfg |
| 2(*) | 53555200 | flash:/system.bin |
| 3(*) | 9959424 | flash:/boot.bin |
| 4 | 3678 | flash:/startup.cfg_backup |
| 5 | 30033 | flash:/default.mdb |
| 6 | 42424 | flash:/startup.mdb |
| 7 | 18 | flash:/pathfile |
| 8 | 232311 | flash:/logfile/logfile.log |
| 9 | 5981 | flash:/startup.cfg_back |
| 10(*) | 6098 | flash:/startup.cfg |
| 11 | 20 | flash:/snmpboots |

Free space: 1009906637 bytes

The current image is boot.bin

(*)-with main attribute

(b)-with backup attribute

(*b)-with both main and backup attribute

2. Enter the number of the file to delete. For example, enter **1** to select the file **testbackup.cfg**.

Please input the file number to change: 1

3. Enter **Y** at the confirmation prompt.

The file you selected is testbackup.cfg,Delete it? (Y/N):Y

Deleting.....Done!

Changing the attribute of software images

Software image attributes include main (M), backup (B), and none (N). System software and boot software can each have multiple none-attribute images but only one main image and one backup image on the switch. You can assign both the M and B attributes to one image. If the M or B attribute you are assigning has been assigned to another image, the assignment removes the attribute from that image. If the removed attribute is the sole attribute of the image, its attribute changes to N.

For example, the system image **system.bin** has the M attribute and the system image **system-update.bin** has the B attribute. After you assign the M attribute to **system-update.bin**, the attribute of **system-update.bin** changes to M+B and the attribute of **system.bin** changes to N.

To change the attribute of a system or boot image:

1. Enter **2** in the Boot menu.

EXTENDED BOOT MENU

1. Download image to flash
2. Select image to boot
3. Display all files in flash
4. Delete file from flash
5. Restore to factory default configuration
6. Enter BootRom upgrade menu
7. Skip current system configuration
8. Set switch startup mode
0. Reboot

Ctrl+Z: Access EXTENDED ASSISTANT MENU
 Ctrl+F: Format file system
 Ctrl+P: Change authentication for console login
 Ctrl+Y: Change Work Mode
 Ctrl+R: Download image to SDRAM and run

Enter your choice(0-8): 2

2. Enter 1 or 2 at the prompt to set the attribute of a software image. (The following output is based on the option 2. To set the attribute of a configuration file, enter 3.)

1. Set image file
2. Set bin file
3. Set configuration file
0. Return to boot menu

Enter your choice(0-3): 2

| File Number | File Size(bytes) | File Name |
|-------------|------------------|--------------------------|
| 1(*) | 53555200 | flash:/system.bin |
| 2(*) | 9959424 | flash:/boot.bin |
| 3 | 13105152 | flash:/boot-update.bin |
| 4 | 91273216 | flash:/system-update.bin |

Free space: 905848832 bytes
 (*)-with main attribute
 (b)-with backup attribute
 (*b)-with both main and backup attribute
 Note:Select .bin files. One but only one boot image and system image must be included.

3. Enter the number of the file you are working with. For example, enter 3 to select the boot image **boot-update.bin** and enter 4 to select the system image **system-update.bin**.

Enter file No.(Allows multiple selection):3
 Enter another file No.(0-Finish choice):4

4. Enter 0 to finish the selection.

Enter another file No.(0-Finish choice):0
 You have selected:
 flash:/boot-update.bin
 flash:/system-update.bin

5. Enter M or B to change its attribute to main or backup. If you change its attribute to M, the attribute of **boot.bin** changes to none.

Please input the file attribute (Main/Backup) M

This operation may take several minutes. Please wait....
Next time, boot-update.bin will become default boot file!
Next time, system-update.bin will become default boot file!
Set the file attribute success!

Handling software upgrade failures

If a software upgrade fails, the system runs the old software version.

To handle a software upgrade failure:

1. Verify that the software release is compatible with the switch model and the correct file is used.
2. Verify that the software release and the Boot ROM release are compatible. For software and Boot ROM compatibility, see the hardware and software compatibility matrix in the correct release notes.
3. Check the physical ports for a loose or incorrect connection.
4. If you are using the console port for file transfer, check the HyperTerminal settings (including the baud rate and data bits) for any wrong setting.
5. Check the file transfer settings:
 - If XMODEM is used, you must set the same baud rate for the terminal as for the console port.
 - If TFTP is used, you must enter the same server IP addresses, file name, and working directory as set on the TFTP server.
 - If FTP is used, you must enter the same FTP server IP address, source file name, working directory, and FTP username and password as set on the FTP server.
6. Check the FTP or TFTP server for any incorrect setting.
7. Check that the storage device has sufficient space for the upgrade file.



Hewlett Packard
Enterprise

HPE 5940-CMW710-R6710P03 Release Notes

Software Feature Changes

The information in this document is subject to change without notice.

© Copyright 2023 Hewlett Packard Enterprise Development LP

Contents

| | |
|---|---|
| About software feature changes | 1 |
| Version R6710P03..... | 1 |
| New feature: DRNI configuraton | 1 |
| New feature: Generating a log message to display product version numbers before and after a software upgrade | 1 |
| Generating a log message to display product version numbers before and after a software upgrade | 1 |
| Log message reference | 2 |
| PKG_UPGRADE_INFO | 2 |
| Modified feature: Enabling SNMP notification for FC interface status changes | 2 |
| Feature change description..... | 2 |
| Command changes | 2 |
| Modified command: enable snmp trap updown | 2 |
| Modified feature: Adding a member to a zone alias | 3 |
| Feature change description..... | 3 |
| Command changes | 3 |
| Modified command: member (zone alias view)..... | 3 |
| Modified feature: Adding a member to a zone | 3 |
| Feature change description..... | 3 |
| Command changes | 3 |
| Modified command: member (zone view) | 3 |
| Modified feature: Displaying parent information for a zone member | 4 |
| Feature change description..... | 4 |
| Command changes | 4 |
| Modified command: display zone member..... | 4 |

| | |
|--|----|
| Release 6710P01 | 5 |
| Release 6710 | 6 |
| New features: Fundamentals features | 7 |
| New features: Layer 2—LAN switching features | 8 |
| New features: Layer 3—IP services features | 11 |
| New features: Layer 3—IP routing features | 15 |
| New features: IP multicast features | 23 |
| New features: MPLS features | 25 |
| New features: ACL and QoS features | 27 |
| New features: Security features | 29 |
| New features: High availability features | 40 |
| New features: Network management and monitoring features | 41 |
| New features: Telemetry features | 44 |
| New features: OpenFlow features | 45 |
| New features: VXLAN features | 46 |
| New features: EVPN features | 46 |
| New features: Intelligent lossless network features | 49 |
| New features: M-LAG support for DRNI commands | 49 |

| | |
|--|-----------|
| New feature: EVPN instance created in system view..... | 50 |
| Configuring an EVPN instance created in system view | 50 |
| Command reference | 51 |
| New command: evpn encapsulation vxlan binding instance..... | 51 |
| New command: evpn instance..... | 51 |
| Modified command: arp mac-learning disable | 52 |
| Modified command: arp-advertising disable..... | 52 |
| Modified command: mac-advertising disable | 53 |
| Modified command: nd-advertising enable | 53 |
| Modified command: nd mac-learning disable..... | 53 |
| Modified command: route-distinguisher | 54 |
| Modified command: vpn-target..... | 54 |
| Modified feature: IPv6 routes with prefixes longer than 64 bits | 55 |
| Feature change description..... | 55 |
| Command changes | 55 |
| Modified command: hardware-resource routing-mode | 55 |
| Modified feature: Match criteria in a traffic class | 55 |
| Feature change description..... | 55 |
| Modified command: if-match | 55 |
| Modified feature: Associating a traffic behavior with a traffic class..... | 56 |
| Feature change description..... | 56 |
| Command changes | 56 |
| Modified command: classifier behavior | 56 |
| Modified feature: Displaying the running configuration..... | 57 |
| Feature change description..... | 57 |
| Command changes | 57 |
| Modified command: display current-configuration..... | 57 |
| Modified feature: Displaying the contents of the configuration file for the next system startup | 58 |
| Feature change description..... | 58 |
| Command changes | 58 |
| Modified command: display saved-configuration | 58 |

| | |
|--|----|
| Modified feature: Optimized display of BGP BMP server information..... | 58 |
| Feature change description..... | 58 |
| Command changes..... | 59 |
| Modified command: display bgp bmp server..... | 59 |
| Modified feature: Disabling BGP session establishment with peers and peer groups | 60 |
| Feature change description..... | 60 |
| Command changes..... | 60 |
| Modified command: ignore all-peers | 60 |
| Modified command: interface-peer/peer ignore | 60 |
| Modified feature: Optimizations to VXLAN command output..... | 61 |
| Feature change description..... | 61 |
| Command changes..... | 61 |
| Modified command: display l2vpn vsi..... | 61 |
| Modified command: display vxlan tunnel | 61 |
| Modified feature: Sharing of VSI interfaces among VSIs..... | 62 |
| Modified feature: Enabling L2TP for the specified protocol | 62 |
| Feature change description..... | 62 |
| Command changes..... | 62 |
| Modified command: l2protocol tunnel dot1q | 62 |
| Modified feature: Creating a service loopback group and specify its service type..... | 63 |
| Feature change description..... | 63 |
| Command changes..... | 63 |
| Modified command: service-loopback group..... | 63 |
| Modified feature: Displaying LSP information | 63 |
| Feature change description..... | 63 |
| Command changes..... | 63 |
| Modified command: display mpls lsp | 63 |

| | |
|--|----|
| Modified feature: Configuring the signaling protocol for an MPLS TE tunnel | |
| interface..... | 64 |
| Feature change description..... | 64 |
| Command changes..... | 64 |
| Modified command: mpls te signaling..... | 64 |
| Modified feature: Configuring the tunnel selection order and set the number of | |
| tunnels for load balancing..... | 64 |
| Feature change description..... | 64 |
| Command changes..... | 65 |
| Modified command: mpls te signaling..... | 65 |
| Modified feature: Verifying MPLS LSP connectivity for an IPv4 prefix..... | 65 |
| Feature change description..... | 65 |
| Command changes..... | 65 |
| Modified command: ping mpls ipv4..... | 65 |
| Modified feature: Tracing MPLS LSPs from the ingress node to the egress | |
| node for an IPv4 prefix..... | 66 |
| Feature change description..... | 66 |
| Command changes..... | 66 |
| Modified command: tracert mpls ipv4..... | 66 |
| Modified feature: Creating a local site..... | 66 |
| Feature change description..... | 66 |
| Command changes..... | 67 |
| Modified command: site..... | 67 |
| Modified feature: Enabling link flapping protection on an interface..... | 67 |
| Feature change description..... | 67 |
| Command changes..... | 67 |
| Modified command: port link-flap protect enable..... | 67 |
| Modified feature: AAA methods in an ISP domain | 68 |
| Feature change description..... | 68 |
| Command reference | 68 |

| | |
|---|-----------|
| Modified command: accounting default..... | 68 |
| Modified command: accounting lan-access | 69 |
| Modified command: accounting login..... | 69 |
| Modified command: accounting portal | 70 |
| Modified command: authentication default..... | 71 |
| Modified command: authentication lan-access | 72 |
| Modified command: authentication login..... | 72 |
| Modified command: authentication portal | 73 |
| Modified command: authorization default | 74 |
| Modified command: authorization lan-access | 74 |
| Modified command: authorization login..... | 75 |
| Modified command: authorization portal | 76 |
| Modified feature: Setting the 802.1X periodic reauthentication timer | 76 |
| Feature change description..... | 76 |
| Command changes | 76 |
| Modified command: dot1x timer | 76 |
| Modified command: dot1x timer reauth-period (interface view) | 77 |
| Modified feature: Setting the periodic MAC reauthentication timer..... | 77 |
| Feature change description..... | 77 |
| Command changes | 77 |
| Modified command: mac-authentication timer (interface view) | 77 |
| Modified command: mac-authentication timer (system view) | 78 |
| Modified feature: Enabling SNMP notifications for L3VPN | 78 |
| Feature change description..... | 78 |
| Command changes | 78 |
| Modified command: snmp-agent trap enable l3vpn | 78 |
| Modified feature: Displaying BGP VPNv4 routing information | 79 |
| Feature change description..... | 79 |
| Command changes | 79 |
| Modified command: display bgp routing-table vpnv4 | 79 |
| Modified feature: Displaying BGP VPNv6 routing information | 81 |
| Feature change description..... | 81 |
| Command changes | 81 |
| Modified command: display bgp routing-table vpnv6 | 81 |

| | |
|--|----|
| Modified feature: Creating an SNMPv3 user | 83 |
| Feature change description..... | 83 |
| Command changes | 83 |
| Modified command: snmp-agent usm-user v3..... | 83 |
| Modified command: snmp-agent calculate-password | 85 |
| Modified feature: Displaying local public keys | 86 |
| Feature change description..... | 86 |
| Command changes | 87 |
| Modified command: display public-key local public..... | 87 |
| Modified feature: Flow-mirroring traffic to an interface | 87 |
| Feature change description..... | 87 |
| Command changes | 87 |
| Modified command: mirror-to interface | 87 |
| Modified feature: Disabling the device from learning the source MAC addresses of Layer 2 protocol packets | 88 |
| Feature change description..... | 88 |
| Command changes | 88 |
| Modified command: mac-address mac-learning pdu | 88 |
| Modified feature: Specifying multiple VLAN IDs for VLAN name TLVs..... | 89 |
| Command changes | 89 |
| Modified command: lldp tlv-enable..... | 89 |
| Modified feature: Restrictions for an interface to join a Layer 2 aggregation group | 92 |
| Feature change description..... | 92 |
| Command changes | 93 |
| Modified feature: Applying a QoS policy to an interface | 93 |
| Feature change description..... | 93 |
| Command changes | 93 |
| Modified command: qos apply policy (interface view)..... | 93 |

| | |
|---|-----|
| Feature 6705 | 94 |
| New feature: FEC mode | 94 |
| Configuring FEC..... | 94 |
| Command reference | 94 |
| port fec mode | 94 |
| New feature: Displaying ND entry statistics | 95 |
| Displaying ND entry statistics | 95 |
| Command reference | 95 |
| display ipv6 neighbors statistics..... | 95 |
| New feature: User IP address conflict SNMP notifications for ARP..... | 97 |
| Enabling user IP address conflict SNMP notifications for ARP..... | 97 |
| Command reference | 97 |
| Modified command: snmp-agent trap enable arp..... | 97 |
| Modified feature: Restrictions for an interface to join a Layer 2 aggregation | |
| group | 98 |
| Feature change description..... | 98 |
| Command changes | 98 |
| Modified feature: Configuring MAC address borrowing | 98 |
| Feature change description..... | 98 |
| Command changes | 98 |
| Modified command: lldp management-address | 98 |
| Modified command: lldp source-mac vlan..... | 99 |
| Modified feature: Configuring the types of advertisable TLVs on a port | 99 |
| Feature change description..... | 99 |
| Command changes | 100 |
| Modified command: lldp tlv-enable..... | 100 |
| Modified feature: Enabling IGMP proxying on an interface | 105 |
| Feature change description..... | 105 |
| Command changes | 105 |
| Modified command: display igmp group..... | 105 |
| Modified command: display igmp proxy group..... | 106 |

| | |
|--|------------|
| Modified command: igmp proxy enable | 106 |
| Modified feature: DRNI term changes | 106 |
| Feature change description..... | 106 |
| Command changes in DRNI | 107 |
| Modified command: display drni consistency..... | 107 |
| Modified command: display drni consistency-check status..... | 107 |
| Modified command: display drni drcp statistics..... | 108 |
| Modified command: display drni keepalive | 108 |
| Modified command: display drni mad verbose..... | 108 |
| Modified command: display drni role..... | 108 |
| Modified command: display drni summary..... | 108 |
| Modified command: display drni system | 109 |
| Modified command: display drni troubleshooting..... | 109 |
| Modified command: display drni verbose..... | 109 |
| Modified command: display drni virtual-ip | 109 |
| Modified command: drni authentication key..... | 109 |
| Modified command: drni auto-recovery reload-delay | 110 |
| Modified command: drni consistency-check disable..... | 110 |
| Modified command: drni consistency-check mode | 110 |
| Modified command: drni drcp period short..... | 111 |
| Modified command: drni ipp mac-address hold | 111 |
| Modified command: drni keepalive { ip ipv6 } | 111 |
| Modified command: drni keepalive hold-time..... | 111 |
| Modified command: drni keepalive interval..... | 112 |
| Modified command: drni mad default-action | 112 |
| Modified command: drni mad exclude interface..... | 112 |
| Modified command: drni mad exclude logical-interfaces | 113 |
| Modified command: drni mad include interface..... | 113 |
| Modified command: drni mad persistent | 113 |
| Modified command: drni mad restore..... | 113 |
| Modified command: drni restore-delay..... | 114 |
| Modified command: drni role priority | 114 |
| Modified command: drni sequence enable | 114 |
| Modified command: drni standalone enable | 114 |
| Modified command: drni system-mac | 115 |
| Modified command: drni system-number | 115 |
| Modified command: drni system-priority | 115 |
| Modified command: port drni group | 115 |
| Modified command: port drni intra-portal-port | 116 |
| Modified command: port drni ipv6 virtual-ip | 116 |
| Modified command: port drni system-mac | 117 |

| | |
|--|-----|
| Modified command: port drni system-priority | 117 |
| Modified command: port drni virtual-ip | 117 |
| Modified command: reset drni drcp statistics | 118 |
| Modified command: reset drni troubleshooting history | 118 |
| Command changes in Track | 119 |
| Modified command: track drni-mad-status | 119 |
| Command changes in portal | 119 |
| Modified command: portal drni load-sharing-mode | 119 |
| Modified command: portal drni traffic backup | 119 |
| Command changes in Web authentication | 120 |
| Modified command: display web-auth user | 120 |
| Command changes in AAA | 120 |
| Modified command: nas-ip (RADIUS scheme view) | 120 |
| Command changes in 802.1X | 120 |
| Modified command: display dot1x connection | 120 |
| Command changes in MAC authentication | 121 |
| Modified command: display mac-authentication connection | 121 |
| Command changes in port security | 121 |
| Modified command: display port-security access-user | 121 |
| Modified command: display port-security static-user connection | 122 |
| Modified command: port-security drni load-sharing-mode | 122 |
| Command changes in DHCP | 122 |
| Modified command: display dhcp snooping drni-statistics | 122 |
| Modified command: display dhcp snooping drni-status | 123 |
| Modified command: reset dhcp snooping drni-statistics | 123 |
| Command changes in DHCPv6 | 123 |
| Modified command: display ipv6 dhcp snooping drni-statistics | 123 |
| Modified command: display ipv6 dhcp snooping drni-status | 123 |
| Modified command: reset ipv6 dhcp snooping drni-statistics | 124 |
| Command changes in VXLAN | 124 |
| Modified command: display l2vpn drni arp | 124 |
| Modified command: display l2vpn drni arp suppression | 124 |
| Modified command: display l2vpn drni mac-address | 125 |
| Modified command: display l2vpn drni nd | 125 |
| Modified command: display l2vpn drni nd suppression | 125 |
| Modified command: display l2vpn service-instance | 125 |
| Modified command: display l2vpn vsi | 126 |
| Command changes in EVPN | 126 |
| Modified command: display evpn drni synchronized-mac | 126 |
| Modified command: evpn drni group | 126 |
| Modified command: evpn drni local | 126 |
| Modified command: l2vpn drni peer-link ac-match-rule vxlan-mapping | 127 |
| Modified command: drni local | 127 |

| | |
|---|------------|
| Modified command: multicast-vpn vxlan drni local..... | 127 |
| Modified command: nexthop evpn-drni group-address..... | 127 |
| Modified command: source | 128 |
| Modified feature: Using OSPF commands on interfaces..... | 128 |
| Feature change description..... | 128 |
| Command changes | 128 |
| Modified feature: Configuring a static route..... | 129 |
| Feature change description..... | 129 |
| Command changes | 129 |
| Modified command: ip route-static | 129 |
| Modified command: ip route-static vpn-instance..... | 131 |
| Modified feature: Configuring an IPv6 static route..... | 133 |
| Feature change description..... | 133 |
| Command changes | 133 |
| Modified command: ipv6 route-static | 133 |
| ESS 6702 | 136 |
| New feature: Specifying a security enhanced level | 136 |
| Specifying a security enhanced level | 136 |
| Command reference | 136 |
| New feature: Private VSI..... | 137 |
| Configuring private VSI | 137 |
| About private VSI | 137 |
| Restrictions and guidelines | 138 |
| Configuring a primary VSI | 139 |
| Configuring a secondary VSI | 139 |
| Command reference | 139 |
| private-vsi..... | 139 |
| private-vsi secondary | 140 |
| Modified feature: Configuring the global priority trust mode for VXLAN packets | |
| | 141 |
| Feature change description..... | 141 |
| Command changes | 141 |

| | |
|--|------------|
| Modified command: qos trust (system view) | 141 |
| Modified feature: Configuring MAC authentication | 141 |
| Feature change description | 141 |
| Command changes | 142 |
| Modified feature: Disabling BGP from flushing all routes to the routing table | |
| | 142 |
| Feature change description | 142 |
| Command changes | 142 |
| Modified command: routing-table bgp-rib-only | 142 |
| Modified feature: Support for HMAC-SHA-256 OSPF/OSPFv3 authentication | |
| mode | 143 |
| Feature change description | 143 |
| Command changes | 143 |
| Modified command: authentication-mode | 143 |
| Modified command: ospf authentication-mode | 143 |
| Modified command: vlink-peer | 144 |
| Modified command: authentication-mode | 144 |
| Modified command: ospfv3 authentication-mode | 145 |
| Modified command: vlink-peer | 145 |
| Modified command: sham-link (OSPF area view) | 146 |
| Modified command: sham-link (OSPFv3 area view) | 147 |
| Modified feature: Displaying the hash keys used for link aggregation load | |
| sharing | 147 |
| Feature change description | 147 |
| Command changes | 147 |
| Modified command: display link-aggregation load-sharing mode | 147 |
| Release 2717 | 149 |
| New feature: Interface alarm functions | 151 |
| Configuring interface alarm functions | 151 |
| Command reference | 153 |
| ifmonitor crc-error | 153 |

| | |
|--|-----|
| ifmonitor input-error..... | 155 |
| ifmonitor input-usage..... | 156 |
| ifmonitor output-error..... | 157 |
| ifmonitor output-usage | 158 |
| port ifmonitor crc-error..... | 159 |
| port ifmonitor input-error..... | 160 |
| port ifmonitor input-usage | 161 |
| port ifmonitor output-error | 162 |
| port ifmonitor output-usage | 163 |
| snmp-agent trap enable ifmonitor | 164 |

New feature: Establishing neighbors through the secondary IP address of an interface..... 165

| | |
|---|-----|
| Enabling OSPF to establish neighbors through the secondary IP address of an interface | 165 |
| Command reference | 166 |
| New command: ospf peer sub-address enable..... | 166 |
| Modified command: display ospf interface | 167 |

New feature: Advertising only the global unicast address in the NEXT_HOP attribute 167

| | |
|--|-----|
| Advertising only the global unicast address in the NEXT_HOP attribute..... | 167 |
| Command reference | 168 |
| nexthop global-address-only | 168 |

New feature: Assigning a VXLAN tunnel interface to a VXLAN tunnel group 169

| | |
|--|-----|
| Assigning a VXLAN tunnel interface to a VXLAN tunnel group | 169 |
| Command reference | 170 |
| group | 170 |

New feature: Setting VLAN tags for sent protocol packets on an interface 170

| | |
|---|-----|
| Enabling the function of setting VLAN tags for sent protocol packets on an interface..... | 170 |
| Command reference | 171 |
| port outbound-vlan-tag enable | 171 |

New feature: Configuring the system ID..... 172

| | |
|--------------------------------|-----|
| Configuring the system ID..... | 172 |
| Command reference | 172 |

| | |
|---|------------|
| sysid | 172 |
| New feature: IPv6 duplicate detection on duplicate addresses | 173 |
| Enabling duplicate detection on duplicate addresses | 173 |
| Command reference | 173 |
| ipv6 address duplicate-detect enable..... | 173 |
| ipv6 address duplicate-detect interval..... | 174 |
| New feature: BGP route re-origination | 175 |
| Configuring BGP route re-origination | 175 |
| Command reference | 176 |
| advertise route-reoriginate | 176 |
| peer advertise vpn-reoriginate ibgp..... | 177 |
| New feature: Sending RA messages over VXLAN tunnels | 178 |
| Enabling a distributed EVPN gateway to send RA messages over VXLAN tunnels..... | 178 |
| Command reference | 178 |
| ipv6 nd ra tunnel-broadcast global enable | 178 |
| New feature: SAVA..... | 179 |
| About SAVA | 179 |
| Benefits | 179 |
| Mechanism | 180 |
| Application scenarios | 180 |
| SAVA tasks at a glance..... | 181 |
| Enabling SAVA..... | 182 |
| Enabling SAVA entry creation based on synchronized remote routes | 182 |
| Adding an interface to a SAVA access group | 183 |
| Configuring SAVA logging..... | 183 |
| Display and maintenance commands for SAVA | 183 |
| Command reference | 184 |
| display ipv6 sava..... | 184 |
| display ipv6 sava packet-drop statistics..... | 185 |
| ipv6 sava access-group | 185 |
| ipv6 sava enable | 186 |
| ipv6 sava log enable spoofing-packet..... | 187 |
| ipv6 sava import remote-route-tag | 188 |
| reset ipv6 sava packet-drop statistics | 188 |
| New feature: IPv6 IS-IS link tag | 189 |
| Configuring the IPv6 IS-IS link tag feature..... | 189 |

| | |
|--|------------|
| Command reference | 190 |
| link-tag inherit enable | 190 |
| isis ipv6 link-tag | 191 |
| New feature: OSPFv3 link tag..... | 191 |
| Configuring the OSPFv3 link tag feature | 191 |
| Command reference | 192 |
| link-tag inherit enable | 192 |
| ospfv3 link-tag | 193 |
| New feature: Displaying a message on enablement of the FTP or Telnet server | |
| | 194 |
| Displaying a message on enablement of the FTP or Telnet server | 194 |
| Command reference | 194 |
| ftp server enable | 194 |
| telnet server enable..... | 195 |
| New feature: Enabling OSPF to advertise the maximum link cost to neighbors | |
| | 195 |
| Enabling OSPF to advertise the maximum link cost to neighbors | 195 |
| About this task..... | 195 |
| Procedure..... | 196 |
| Command reference | 196 |
| ospf peer hold-max-cost duration..... | 196 |
| New feature: Enabling IS-IS to advertise the maximum link cost to neighbors | |
| | 197 |
| Enabling IS-IS to advertise the maximum link cost to neighbors | 197 |
| About this task..... | 197 |
| Procedure..... | 197 |
| Command reference | 197 |
| isis peer hold-max-cost duration | 197 |
| New feature: Enabling OSPFv3 to advertise the maximum link cost to | |
| neighbors..... | 198 |
| Enabling OSPFv3 to advertise the maximum link cost to neighbors | 198 |
| About this task..... | 198 |

| | |
|---|-----|
| Procedure..... | 198 |
| Command reference | 199 |
| ospfv3 peer hold-max-cost duration..... | 199 |

New feature: Enabling BGP to advertise a default route to a peer or peer group

| | |
|-------|-----|
| | 200 |
|-------|-----|

| | |
|---|-----|
| Enabling BGP to advertise a default route to a peer or peer group | 200 |
| Command reference | 200 |
| Modified command: peer default-route-advertise..... | 200 |

New feature: SAVI 202

| | |
|--|-----|
| Configuring SAVI..... | 202 |
| About SAVI..... | 202 |
| SAVI application scenarios | 202 |
| SAVI tasks at a glance | 203 |
| Enabling SAVI | 203 |
| Configuring IPv6 source guard..... | 203 |
| Configuring DHCPv6 snooping | 203 |
| Configuring ND parameters | 203 |
| Setting the entry deletion delay..... | 204 |
| Enabling packet spoofing logging and filtering entry logging | 204 |
| SAVI configuration examples..... | 205 |
| Example: Configuring DHCPv6-only SAVI..... | 205 |
| Example: Configuring SLAAC-only SAVI..... | 206 |
| Example: Configuring DHCPv6+SLAAC SAVI..... | 207 |
| Command reference | 209 |
| ipv6 savi down-delay..... | 209 |
| ipv6 savi log enable..... | 209 |
| ipv6 savi strict..... | 210 |

New feature: Setting the maximum number of IPv6SG bindings on an interface

| | |
|-------|-----|
| | 211 |
|-------|-----|

| | |
|--|-----|
| Setting the maximum number of IPv6SG bindings on an interface..... | 211 |
| Command reference | 211 |
| ipv6 verify source max-entries | 211 |

| | |
|---|------------|
| New feature: Setting the interval at which the SNMP module examines the system configuration for changes | 212 |
| Setting the interval at which the SNMP module examines the system configuration for changes..... | 212 |
| Command reference | 212 |
| snmp-agent configuration-examine interval | 212 |
| New feature: Displaying SNMP notifications drop records | 213 |
| Displaying SNMP notifications drop records | 213 |
| Command reference | 213 |
| display snmp-agent trapbuffer drop | 213 |
| New feature: Displaying SNMP notifications sending records | 214 |
| Displaying SNMP notifications sending records..... | 214 |
| Command reference | 214 |
| display snmp-agent trapbuffer send..... | 214 |
| New feature: Clearing all records from the SNMP trap buffer | 215 |
| Clearing all records from the SNMP trap buffer | 215 |
| Command reference | 215 |
| reset snmp-agent trapbuffer..... | 215 |
| New feature: PWDCTL messages | 215 |
| PWDCTL_UPDATETIME..... | 215 |
| New feature: LIPC messages | 216 |
| PORT_CHANGE..... | 216 |
| New feature: Applying an SSL server policy to the NETCONF over SOAP over HTTPS service | 216 |
| Applying an SSL server policy to the NETCONF over SOAP over HTTPS service..... | 216 |
| Command reference | 216 |
| netconf soap https ssl-server-policy..... | 216 |
| New feature: Policing ND protocol packet to the control plane..... | 217 |
| Configuring ND protocol packet policing..... | 217 |
| Command reference | 218 |
| Modified command: display qos policy control-plane pre-defined..... | 218 |

| | |
|--|------------|
| Modified command: if-match | 218 |
| New feature: RA message transmission over VXLAN tunnels between distributed EVPN gateways | 219 |
| Enabling a distributed EVPN gateway to send RA messages over VXLAN tunnels..... | 219 |
| Command reference | 219 |
| ipv6 nd ra tunnel-broadcast enable..... | 219 |
| New feature: MAC move logging for VSIs..... | 220 |
| Displaying MAC move records for VSIs..... | 220 |
| Command reference | 220 |
| display l2vpn mac-address mac-move | 220 |
| reset l2vpn mac-address mac-move | 221 |
| New feature: MAC learning limits for VSIs | 221 |
| Setting the MAC learning limit on a VSI..... | 221 |
| Command reference | 222 |
| mac-table limit | 222 |
| mac-table limit drop-unknown | 222 |
| New feature: DRNI in EVPN-DCI | 223 |
| Configuring DRNI in EVPN-DCI | 223 |
| Command reference | 224 |
| nexthop evpn-drni group-address | 224 |
| New feature: L2VPN statistics collection interval | 225 |
| Setting the L2VPN statistics collection interval | 225 |
| Command reference | 225 |
| l2vpn statistics interval | 225 |
| New feature: Unsolicited NA learning | 226 |
| Enabling unsolicited NA learning | 226 |
| Command reference | 226 |
| ipv6 nd unsolicited-na-learning enable..... | 226 |
| New feature: IPv6 enhancement for OpenFlow flow entries..... | 227 |
| Enabling IPv6 enhancement for OpenFlow flow entries | 227 |
| Command reference | 227 |
| openflow flow-table ipv6-enhanced..... | 227 |

| | |
|---|------------|
| New feature: IPv6-matching QoS policy..... | 228 |
| Configuring an IPv6-matching QoS policy | 228 |
| Command reference | 229 |
| Modified command: qos policy..... | 229 |
| Modified command: qos apply policy | 230 |
| Modified command: qos apply policy global..... | 230 |
| Modified command: display qos policy..... | 231 |
| Modified command: display qos policy interface..... | 231 |
| Modified command: display qos policy global | 232 |
| Modified command: reset qos policy global | 232 |
| New feature: Enabling BGP to not flush specific routes to the routing table | 233 |
| Enabling BGP to not flush specific routes to the routing table | 233 |
| Command reference | 234 |
| routing-table bgp-rib-only | 234 |
| New feature: Enabling BGP to reset peer sessions gracefully | 235 |
| Enabling BGP to reset peer sessions gracefully..... | 235 |
| Command reference | 235 |
| graceful-restart peer-reset..... | 235 |
| New feature: Setting the extra time to wait after the restart timer expires .. | 236 |
| Setting the extra time to wait after the restart timer expires..... | 236 |
| Command reference | 236 |
| peer graceful-restart timer restart extra..... | 236 |
| New feature: Enabling password change prompt logging | 237 |
| Enabling password change prompt logging | 237 |
| Command reference | 238 |
| local-server log change-password-prompt | 238 |
| New feature: Enabling mandatory weak password change..... | 239 |
| Enabling mandatory weak password change | 239 |
| Command reference | 239 |
| New command: password-control change-password weak-password enable..... | 239 |
| Modified command: display password-control | 240 |
| New feature: Configuring ND packet rate limit | 241 |
| Configuring ND packet rate limit | 241 |

| | |
|--|------------|
| Command reference | 242 |
| ipv6 nd rate-limit | 242 |
| ipv6 nd rate-limit log enable | 243 |
| ipv6 nd rate-limit log interval | 243 |
| New feature: gRPC logging in dial-in mode | 244 |
| Enabling gRPC logging in dial-in mode | 244 |
| Command reference | 244 |
| grpc log dial-in gnmi | 244 |
| grpc log dial-in rpc | 245 |
| New feature: gRPC logging in dial-out mode | 246 |
| Enabling gRPC logging in dial-out mode | 246 |
| Command reference | 246 |
| grpc log dial-out | 246 |
| New feature: Creating frame match criteria based on VXLAN IDs for dynamic ACs on the IPL of an EVPN DR system | 247 |
| Creating frame match criteria based on VXLAN IDs for dynamic ACs on the IPL of an EVPN DR system .. | 247 |
| Command reference | 248 |
| l2vpn drni peer-link ac-match-rule vxlan-mapping | 248 |
| New feature: Removing Layer 3 VPN route targets from BGP EVPN routes for VXLAN mapping | 249 |
| Removing Layer 3 VPN route targets from BGP EVPN routes for VXLAN mapping | 249 |
| Command reference | 249 |
| mapping-vni remove vpn-target | 249 |
| New feature: Enabling asymmetric IRB mode for EVPN VXLAN traffic forwarding | 250 |
| Configuring the traffic forwarding mode for EVPN VXLAN | 250 |
| Command reference | 251 |
| evpn irb asymmetric | 251 |

| | |
|--|-----|
| New feature: Setting the Gateway IP Address field for EVPN IP prefix advertisement routes | 251 |
| Setting the Gateway IP Address field for EVPN IP prefix advertisement routes | 251 |
| Command reference | 252 |
| apply gateway-ip | 252 |
| apply ipv6 gateway-ip | 252 |
| New feature: Enabling BGP to redistribute all IGP routes with the same destination address | 253 |
| Enabling BGP to redistribute all IGP routes with the same destination address | 253 |
| Command reference | 254 |
| import-route multipath | 254 |
| New feature: Policing ND protocol packets | 255 |
| Policing ND protocol packets | 255 |
| Command reference | 255 |
| Modified command: display qos policy control-plane pre-defined | 255 |
| Modified command: if-match | 255 |
| New feature: Using an ACL to control TCP connections from IPv4 HTTP or HTTPS clients | 256 |
| Using an ACL to control TCP connections from IPv4 HTTP or HTTPS clients | 256 |
| Command reference | 256 |
| http acl | 256 |
| https acl | 257 |
| New feature: Enabling Layer 3 packet statistics counting | 258 |
| Enabling Layer 3 packet statistics counting | 258 |
| Command reference | 258 |
| statistics l3-packet enable | 258 |
| New feature: Specifying the PKI domain for secure communication with collectors | 259 |
| Command reference | 260 |
| grpc pki domain | 260 |

| | |
|--|-----|
| New feature: DRNI MAD | 260 |
| Configuring DRNI MAD | 260 |
| About this task | 260 |
| Configuring the default DRNI MAD action on network interfaces | 261 |
| Specifying interfaces to be shut down by DRNI MAD when the DR system splits | 262 |
| Command reference | 262 |
| New command: drni mad default-action | 262 |
| New command: drni mad include interface | 264 |
| Modified command: display drni mad verbose | 265 |
| New feature: Configuring support for forwarding DHCPv6 replies based on MAC address table | 267 |
| Configuring support for forwarding DHCPv6 replies based on MAC address table | 267 |
| Command reference | 267 |
| ipv6 dhcp relay mac-forward enable | 267 |
| New feature: Configuring basic DHCPv6 snooping features on a VXLAN network | 268 |
| Configuring basic DHCPv6 snooping features on a VXLAN network | 268 |
| Command changes | 269 |
| ipv6 dhcp snooping binding record | 269 |
| ipv6 dhcp snooping trust tunnel | 270 |
| Modified command: ipv6 dhcp snooping trust | 271 |
| New feature: Configuring BGP VPNv6 routes to use private network next hops | 271 |
| Configuring BGP VPNv6 routes to use private network next hops | 271 |
| Command reference | 272 |
| peer next-hop-vpn | 272 |
| New feature: Configuring BGP route replication between public and VPN instances | 273 |
| Configuring BGP route replication between public and VPN instances | 273 |
| Command reference | 274 |
| route-replicate enable | 274 |
| export route-policy | 275 |

| | |
|---|------------|
| import route-policy..... | 275 |
| New feature: Displaying IPv6 multicast group information that MLD snooping learns from EVPN..... | 276 |
| Displaying IPv6 multicast group information that MLD snooping learns from EVPN | 276 |
| Command reference | 277 |
| display mld-snooping evpn-group | 277 |
| New feature: Enabling password control blacklist for all user line types..... | 277 |
| Enabling password control blacklist for all user line types | 277 |
| Command changes | 278 |
| password-control blacklist all-line..... | 278 |
| New feature: Configuring ND attack detection for a VSI..... | 279 |
| Configuring ND attack detection for a VSI | 279 |
| Command reference | 280 |
| Modified command: ipv6 nd detection enable..... | 280 |
| Modified command: ipv6 nd detection trust..... | 280 |
| Modified command: display ipv6 nd detection statistics..... | 280 |
| Modified command: reset ipv6 nd detection statistics..... | 281 |
| New feature: MVXLAN..... | 281 |
| MVXLAN overview | 281 |
| Restrictions: Underlay network restriction..... | 281 |
| MVXLAN benefits..... | 281 |
| MVXLAN network model..... | 281 |
| Basic concepts in MVXLAN | 282 |
| How MVXLAN works..... | 282 |
| MDT-based transmission | 282 |
| MP-BGP extension for MVXLAN..... | 283 |
| Automatic MVXLAN tunnel establishment and assignment..... | 283 |
| Default MDT establishment..... | 283 |
| Default MDT-based transmission..... | 284 |
| MDT switchover | 284 |
| MVXLAN tasks at a glance | 285 |
| Enabling IP multicast routing for a VPN instance..... | 286 |
| Creating an MVXLAN..... | 286 |
| Configuring a default group..... | 287 |
| Specifying the MVXLAN source interface | 287 |
| Configuring MDT switchover parameters..... | 288 |

| | |
|--|-----|
| Configuring a VSI interface as a distributed DR interface..... | 288 |
| Configuring an MVXLAN extranet RPF selection policy | 289 |
| Display and maintenance commands for MVXLAN | 289 |
| MVXLAN configuration examples | 290 |
| Example: Configuring inter-MVXLAN Layer 3 multicast forwarding..... | 290 |
| Example: Configuring MVXLAN Layer 3 multicast forwarding between private and public networks.... | 301 |
| MVXLAN commands..... | 314 |
| address-family ipv4 | 314 |
| data-delay | 314 |
| data-group..... | 315 |
| default-group | 316 |
| display multicast-vpn vxlan data-group receive | 317 |
| display multicast-vpn vxlan data-group send | 318 |
| display multicast-vpn vxlan default-group | 320 |
| multicast extranet select-rpf | 321 |
| multicast-vpn vxlan public-instance..... | 322 |
| multicast-vpn vxlan vpn-instance | 323 |
| pim distributed-dr | 324 |
| source | 324 |

Modified feature: Configuring LLDP to advertise IPv4 and IPv6 management

addresses at the same time..... 325

| | |
|---|-----|
| Feature change description..... | 325 |
| Command changes | 325 |
| Modified command: lldp global tlv-enable basic-tlv management-address-tlv | 325 |

Modified feature: Specifying a padding mode for the Interface-ID option on the

DHCPv6 relay agent..... 326

| | |
|---|-----|
| Feature change description..... | 326 |
| Command changes | 326 |
| Modified command: ipv6 dhcp relay interface-id..... | 326 |

Modified feature: Displaying monitor link group information 326

| | |
|---|-----|
| Feature change description..... | 326 |
| Command changes | 327 |
| Modified command: display monitor-link group..... | 327 |

| | |
|---|-----|
| Modified feature: Displaying the path that the IPv6 packets traverse from source to destination..... | 327 |
| Feature change description..... | 327 |
| Command changes..... | 327 |
| Modified command: tracert ipv6..... | 327 |
| Modified feature: Exchanging both IPv4 and IPv6 routes with a BGP peer | 328 |
| Feature change description..... | 328 |
| Command changes..... | 328 |
| Modified feature: Enabling load balancing for BGP routes that have different AS_PATH attributes of the same length | 331 |
| Feature change description..... | 331 |
| Command changes..... | 331 |
| Modified command: balance as-path-relax | 331 |
| Modified feature: Enabling BGP to exchange labeled routes with a peer or peer group | 332 |
| Feature change description..... | 332 |
| Command changes..... | 332 |
| Modified command: peer label-route-capability | 332 |
| Modified feature: Configuring a static neighbor entry | 332 |
| Feature change description..... | 332 |
| Command changes..... | 332 |
| Modified command: ipv6 neighbor | 332 |
| Modified command: display ipv6 neighbors | 334 |
| Modified feature: Modifying information in received BGP EVPN IP prefix advertisement routes on an ED | 334 |
| Feature change description..... | 334 |
| Command changes..... | 334 |
| Modified command: peer re-originated | 334 |

| | |
|---|-----|
| Modified feature: Configuring passive TCP connection settings for a VXLAN | |
| OVSDb VTEP..... | 335 |
| Feature change description..... | 335 |
| Command changes..... | 335 |
| Modified command: ovsdb server ptcp | 335 |
| Modified feature: Setting the password for local password authentication · | 336 |
| Feature change description..... | 336 |
| Command changes..... | 336 |
| Modified command: set authentication password | 336 |
| Modified feature: Enabling an ED to replace the L3 VXLAN ID, RD, and route | |
| targets of BGP EVPN routes..... | 336 |
| Enabling an ED to replace the L3 VXLAN ID, RD, and route targets of BGP EVPN routes | 336 |
| Command reference | 337 |
| Modified command: peer re-originated | 337 |
| New command: peer advertise original-route | 338 |
| New command: peer suppress re-originated | 339 |
| Modified feature: Displaying Ethernet service instance information | 340 |
| Feature change description..... | 340 |
| Command changes..... | 340 |
| Modified command: display l2vpn service-instance | 340 |
| Modified feature: Configuring OSPF FRR | 341 |
| Feature change description..... | 341 |
| Command changes..... | 341 |
| Modified command: fast-reroute | 341 |
| Modified feature: Displaying routing table information..... | 341 |
| Feature change description..... | 341 |
| Command changes..... | 341 |
| Modified command: display ip routing-table..... | 341 |

| | |
|--|-----|
| Modified feature: Configuring the maximum number of Add-Path optimal routes that can be advertised to all peers | 342 |
| Feature change description..... | 342 |
| Command changes..... | 342 |
| Modified command: additional-paths select-best..... | 342 |
| Modified feature: Configuring the BGP Additional Paths capabilities | 343 |
| Feature change description..... | 343 |
| Command changes..... | 343 |
| Modified command: peer additional-paths | 343 |
| Modified feature: Configuring the maximum number of Add-Path optimal routes that can be advertised to a peer or peer group | 344 |
| Feature change description..... | 344 |
| Command changes..... | 344 |
| Modified command: peer advertise additional-paths best..... | 344 |
| Modified feature: Password control restriction features..... | 345 |
| Feature change description..... | 345 |
| Command changes..... | 346 |
| Modified command: password-control composition | 346 |
| Modified command: password-control super composition | 346 |
| Modified command: password-control complexity | 346 |
| Modified feature: Flow-mirroring traffic to interfaces | 347 |
| Feature change description..... | 347 |
| Command changes..... | 347 |
| Modified command: mirror-to interface | 347 |
| Modified feature: Configuring ARP scanning | 348 |
| Feature change description..... | 348 |
| Command changes..... | 348 |
| Modified command: arp scan | 348 |
| Modified feature: NETCONF logging | 349 |
| Feature change description..... | 349 |
| Command changes..... | 349 |

| | |
|--|-----|
| Modified command: netconf log | 349 |
| Modified feature: Enabling route router MAC replacement for a peer or peer group on an EVPN-DCI ED..... | 350 |
| Feature change description..... | 350 |
| Command changes | 350 |
| Modified command: peer router-mac-local..... | 350 |
| Modified feature: Value range change for the extended community attribute list number in an RR reflection policy | 350 |
| Feature change description..... | 350 |
| Command changes | 351 |
| Modified command: rr-filter | 351 |
| Modified feature: Allowing a DR interface to be up when it is the only member of a DR group | 351 |
| Feature change description..... | 351 |
| Command changes | 351 |
| Modified command: port drni group | 351 |
| Modified feature: Configuring a static route..... | 352 |
| Feature change description..... | 352 |
| Command changes | 352 |
| Modified command: ip route-static | 352 |
| Modified command: ip route-static vpn-instance..... | 353 |
| Modified feature: Configuring an IPv6 static route..... | 355 |
| Feature change description..... | 355 |
| Command changes | 355 |
| Modified command: ipv6 route-static | 355 |
| Modified feature: Displaying OSPFv3 LSDB information..... | 357 |
| Feature change description..... | 357 |
| Command changes | 357 |
| Modified command: display ospfv3 lsdb | 357 |

| | |
|--|-----|
| Modified feature: Configuring BGP extended community lists..... | 357 |
| Feature change description..... | 357 |
| Command changes..... | 357 |
| Modified command: display ip extcommunity-list..... | 357 |
| Modified command: if-match extcommunity..... | 358 |
| Modified command: ip extcommunity-list..... | 358 |
| Modified feature: Replicating VLINK direct routes..... | 359 |
| Feature change description..... | 359 |
| Command changes..... | 359 |
| Modified command: route-replicate..... | 359 |
| Modified feature: Specifying a controller for an OpenFlow switch and configuring the main connection to the controller..... | 360 |
| Feature change description..... | 360 |
| Command changes..... | 361 |
| Modified command: controller address..... | 361 |
| Modified feature: Displaying the CPU usage statistics | 361 |
| Feature change description..... | 361 |
| Command changes..... | 361 |
| Modified command: display cpu-usage..... | 361 |
| Modified feature: Displaying memory usage information..... | 362 |
| Feature change description..... | 362 |
| Command changes..... | 362 |
| Modified command: display memory..... | 362 |
| Modified feature: Password complexity checking..... | 362 |
| Feature change description..... | 362 |
| Command changes..... | 363 |
| Modified feature: Configuring kernel thread deadlock detection | 363 |
| Feature change description..... | 363 |
| Command changes..... | 363 |
| Modified command: monitor kernel deadlock enable | 363 |

Modified feature: Configuring a traffic redirecting action in a QoS traffic

behavior..... 363

Feature change description..... 363

Command changes..... 363

Modified command: redirect..... 363

Modified feature: DRNI IPP configuration 364

Feature change description..... 364

Command changes..... 364

| | |
|--|-----|
| Release 2702 | 365 |
| New features: Fundamentals features | 366 |
| New features: Layer 2—LAN switching features | 366 |
| New features: Layer 3—IP services features | 367 |
| New features: Layer 3—IP routing features | 367 |
| New features: IP multicast features | 370 |
| New features: MPLS features | 370 |
| New features: ACL and QoS features | 371 |
| New features: Security features | 371 |
| New features: High availability features | 372 |
| New features: OpenFlow features | 372 |
| New features: VXLAN features | 373 |
| New features: EVPN features | 373 |
| Modified feature: Software patching by using issu commands | 373 |
| Feature change description | 373 |
| Command changes | 374 |
| Modified command: display version comp-matrix | 374 |
| Modified command: issu load | 374 |
| Modified command: issu one-step | 375 |
| Modified feature: Automatic configuration | 375 |
| Feature change description | 375 |
| Command changes | 375 |

| | |
|---|-----|
| Modified feature: Collision handling in BFD MAD..... | 376 |
| Feature change description..... | 376 |
| Command changes..... | 376 |
| Modified feature: Processing after the link mode of an Ethernet interface is switched | 376 |
| Feature change description..... | 376 |
| Command changes..... | 376 |
| Modified feature: Link state change suppression on an interface..... | 376 |
| Feature change description..... | 376 |
| Command changes..... | 377 |
| Modified command: link-delay..... | 377 |
| Modified feature: Setting the global aging timer for ND entries in stale state | 377 |
| Feature change description..... | 377 |
| Command changes..... | 378 |
| Modified command: ipv6 neighbor stale-aging..... | 378 |
| Modified feature: Setting the interface-specific aging timer for ND entries in stale state | 378 |
| Feature change description..... | 378 |
| Command changes..... | 378 |
| Modified command: ipv6 neighbor timer stale-aging..... | 378 |
| Modified feature: Configuring OSPF area authentication | 379 |
| Feature change description..... | 379 |
| Command changes..... | 379 |
| Modified command: authentication-mode | 379 |
| Modified feature: Configuring OSPF interface authentication..... | 379 |
| Feature change description..... | 379 |
| Command changes..... | 380 |
| Modified command: ospf authentication-mode | 380 |

| | |
|--|-----|
| Modified feature: Configuring a virtual link | 380 |
| Feature change description..... | 380 |
| Command changes | 381 |
| Modified command: vlink-peer | 381 |
| Modified feature: Displaying statistics for the IGMP messages and PIMv2 hello messages learned through IGMP snooping | 381 |
| Feature change description..... | 381 |
| Command changes | 381 |
| Modified command: display igmp-snooping statistics | 381 |
| Modified feature: Displaying statistics for the MLD messages and IPv6 PIM hello messages learned through MLD snooping | 383 |
| Feature change description..... | 383 |
| Command changes | 384 |
| Modified command: display mld-snooping statistics | 384 |
| Modified feature: Enabling a VPN instance to replicate routes from the public network or other VPN instances..... | 385 |
| Feature change description..... | 385 |
| Command changes | 386 |
| Modified command: route-replicate..... | 386 |
| Modified feature: Using BFD to verify LSP or MPLS TE tunnel connectivity | 386 |
| Feature change description..... | 386 |
| Command changes | 386 |
| Modified command: mpls bfd | 386 |
| Modified feature: MAC authentication VLAN mode | 387 |
| Feature change description..... | 387 |
| Command changes | 387 |
| Modified command: mac-authentication host-mode multi-vlan | 387 |
| Modified feature: Port security MAC move | 388 |
| Feature change description..... | 388 |
| Command changes | 388 |

| | |
|---|------------|
| Modified command: port-security mac-move permit | 388 |
| Modified feature: Creating a BFD session for detecting the local interface state | |
| | 388 |
| Feature change description..... | 388 |
| Command changes | 389 |
| Modified command: bfd detect-interface | 389 |
| Modified feature: Associating Track with application modules..... | 389 |
| Feature change description..... | 389 |
| Command changes | 389 |
| Modified command: track bfd ctrl | 389 |
| Modified command: track bfd echo | 390 |
| Modified command: track cfd | 390 |
| Modified command: track interface | 390 |
| Modified command: track interface physical | 390 |
| Modified command: track interface protocol | 391 |
| Modified command: track ip route reachability..... | 391 |
| Modified command: track lldp neighbor | 391 |
| Modified command: track nqa | 392 |
| Modified feature: Configuring an interface event for a CLI-defined EAA monitor | |
| policy | 392 |
| Feature change description..... | 392 |
| Command changes | 392 |
| Modified command: event-interface | 392 |
| Modified feature: Configuring an EAA monitor policy by using Tcl | 393 |
| Feature change description..... | 393 |
| Command changes | 393 |
| Modified feature: Removing a TCP or UDP listening service for a VPN instance | |
| | 393 |
| Feature change description..... | 393 |
| Command changes | 393 |
| Modified command: nqa server tcp-connect | 393 |
| Modified command: nqa server udp-echo..... | 394 |

| | |
|--|-----|
| Modified feature: Specifying the role of the device in the VCF fabric | 394 |
| Feature change description..... | 394 |
| Command reference | 394 |
| Modified command: vcf-fabric role | 394 |
| Modified command: display vcf-fabric role..... | 395 |
| Modified command: display vcf-fabric underlay autoconfigure | 396 |
| Modified feature: Frame match criteria of VXLAN Ethernet service instances | 399 |
| Feature change description..... | 399 |
| Command changes | 399 |
| Modified command: encapsulation..... | 399 |
| Modified feature: NETCONF logging | 399 |
| Feature change description..... | 399 |
| Command changes | 400 |
| Modified command: netconf log | 400 |
| Feature 2701 | 401 |
| New feature: Filtering the output from a display command by using multiple conditions | 405 |
| Filtering the output from a display command by using multiple conditions | 405 |
| Command reference | 405 |
| display [[[by-linenum] { begin exclude include }]..... | 405 |
| New feature: Executing a batch file..... | 406 |
| Executing a batch file..... | 406 |
| Command reference | 406 |
| execute..... | 406 |
| New feature: Disabling automatic system-wide next-startup configuration file operations..... | 407 |
| Disabling automatic system-wide next-startup configuration file operations..... | 407 |
| Command reference | 408 |
| standby auto-update config..... | 408 |

| | |
|---|-----|
| New feature: Using an encrypted configuration file to roll back configuration | 409 |
| Using an encrypted configuration file to roll back configuration | 409 |
| Command reference | 409 |
| Modified command: configuration replace file | 409 |
| New feature: Support for encrypted configuration files for configuration comparison | 409 |
| Comparing encrypted configuration files with other configurations for their differences | 409 |
| Command reference | 409 |
| Modified command: display current-configuration diff | 409 |
| Modified command: display diff | 410 |
| New feature: ND MAD on management Ethernet ports | 410 |
| Configuring ND MAD on management Ethernet ports | 410 |
| Command reference | 412 |
| mad nd enable | 412 |
| New feature: Fast retrain | 413 |
| Configuring fast retrain | 413 |
| Command reference | 414 |
| port fast-retrain enable | 414 |
| New feature: Displaying the status and packet statistics of interfaces | 414 |
| Displaying the status and packet statistics of interfaces | 414 |
| Command reference | 415 |
| display interface link-info | 415 |
| New feature: Displaying operating status and information of all interfaces except subinterfaces | 416 |
| Displaying operating status and information of all interfaces except subinterfaces | 416 |
| Command reference | 416 |
| display interface main | 416 |
| New feature: Enabling PFC on all interfaces | 427 |
| Enabling PFC on all interfaces | 427 |

| | |
|---|------------|
| Command reference | 427 |
| priority-flow-control (system view) | 427 |
| Command reference | 428 |
| priority-flow-control no-drop dot1p (system view) | 428 |
| New feature: Configuring automatic link aggregation | 429 |
| Configuring automatic link aggregation | 429 |
| Command reference | 430 |
| link-aggregation auto-aggregation enable | 430 |
| New feature: Disabling the default action of selecting a Selected port for dynamic aggregation groups that have not received LACPDUs | 431 |
| Disabling the default action of selecting a Selected port for dynamic aggregation groups that have not received LACPDUs | 431 |
| Command reference | 431 |
| lacp default-selected-port disable | 431 |
| New feature: Physical state change suppression on aggregate interfaces .. | 432 |
| Configuring physical state change suppression on an aggregate interface | 432 |
| Command reference | 433 |
| link-delay | 433 |
| New feature: VXLAN multiport ARP entries | 434 |
| Configuring VXLAN multiport ARP entries | 434 |
| Command reference | 435 |
| New command: arp multiport | 435 |
| New command: mac-address multiport vsi | 436 |
| New command: mac-address multicast vsi | 438 |
| Modified command: display arp | 440 |
| Modified command: display l2vpn mac-address | 442 |
| New feature: Enabling recording user port move for ARP | 445 |
| Enabling recording user port move for ARP | 445 |
| Display and maintenance commands for ARP | 446 |
| Command reference | 446 |
| arp user-move record enable | 446 |
| display arp user-move record | 446 |

| | |
|---|------------|
| New feature: Enabling recording user IP address conflicts | 448 |
| Enabling recording user IP address conflicts..... | 448 |
| Display and maintenance commands for ARP | 448 |
| Command reference | 448 |
| arp user-ip-conflict record enable..... | 448 |
| display arp user-ip-conflict record | 449 |
| New feature: Enabling interface consistency check for ARP and MAC address entries..... | 451 |
| Enabling interface consistency check for ARP and MAC address entries | 451 |
| Command reference | 451 |
| arp mac-interface-consistency check enable | 451 |
| New feature: Specifying a DHCP relay address pool for DHCP clients | 452 |
| Specifying a DHCP relay address pool for DHCP clients | 452 |
| Command reference | 453 |
| Modified command:remote-server | 453 |
| dhcp relay pool..... | 453 |
| New feature: Configuring the padding mode for the Vendor-Specific sub-option | 455 |
| Configuring the padding mode for the Vendor-Specific sub-option | 455 |
| Command reference | 456 |
| dhcp snooping information vendor-specific..... | 456 |
| New feature: Configuring IPv6 ND direct route advertisement | 457 |
| Configuring IPv6 ND direct route advertisement..... | 457 |
| About IPv6 ND direct route advertisement..... | 457 |
| Layer 3 access network application | 457 |
| Procedure..... | 457 |
| Command reference | 457 |
| ipv6 nd route-direct advertise..... | 457 |
| New feature: Configuring ND snooping in a VXLAN | 458 |
| Configuring ND snooping in a VXLAN | 458 |
| About ND snooping in a VXLAN | 458 |
| Procedure..... | 459 |

| | |
|---|------------|
| Command reference | 459 |
| Modified command: display ipv6 nd snooping count..... | 459 |
| Modified command: ipv6 nd snooping enable global | 459 |
| Modified command: ipv6 nd snooping enable link-local..... | 460 |
| New command: display ipv6 nd snooping count vsi..... | 460 |
| New command: display ipv6 nd snooping vsi | 461 |
| New command: reset ipv6 nd snooping vsi..... | 462 |
| New feature: Enabling recording user IPv6 address conflicts | 462 |
| Enabling recording user IPv6 address conflicts | 462 |
| Command reference | 463 |
| ipv6 nd user-ip-conflict record enable | 463 |
| display ipv6 nd user-ip-conflict record..... | 463 |
| New feature: Enabling recording user port move for ND | 465 |
| Enabling recording user port move for ND..... | 465 |
| Command reference | 465 |
| ipv6 nd user-move record enable..... | 465 |
| display ipv6 nd user-move record | 466 |
| New feature: Enabling ND logging for user online and offline events | 467 |
| Enabling ND logging for user online and offline events | 467 |
| Command reference | 468 |
| ipv6 nd online-offline-log enable | 468 |
| New feature: Specifying the boot file URL in RA messages | 468 |
| Specifying the boot file URL in RA messages | 468 |
| Command reference | 469 |
| ipv6 nd ra boot-file-url | 469 |
| New feature: Setting the aging timer for ND entries in stale state on an interface | |
| | 469 |
| Setting the aging timer for ND entries in stale state on an interface | 469 |
| Command reference | 470 |
| ipv6 neighbor timer stale-aging..... | 470 |
| New feature: Specifying DNS server information in RA messages | 471 |
| Specifying DNS server information in RA messages | 471 |
| Command reference | 471 |

| | |
|--|------------|
| ipv6 nd ra dns server..... | 471 |
| New feature: Specifying the DNS domain name suffix information in RA | |
| messages | 473 |
| Specifying the DNS domain name suffix information in RA messages..... | 473 |
| Command reference | 473 |
| ipv6 nd ra dns search-list | 473 |
| New feature: Suppressing DNS information in RA messages..... | 475 |
| Suppressing DNS information in RA messages..... | 475 |
| Command reference | 476 |
| ipv6 nd ra dns server suppress | 476 |
| ipv6 nd ra dns search-list suppress..... | 477 |
| New feature: Enabling the DHCPv6 relay agent to advertise host routes for | |
| assigned IPv6 addresses..... | 478 |
| Enabling the DHCPv6 relay agent to advertise host routes for assigned IPv6 addresses | 478 |
| Command reference | 478 |
| New feature: Enabling IPv6 address binding conversion for IP source guard | |
| | 479 |
| Enabling IPv6 address binding conversion for IP source guard | 479 |
| Command reference | 480 |
| ipv6 dhcp server entry-convert enable | 480 |
| New feature: Configuring DHCPv6 guard | 481 |
| About DHCPv6 guard..... | 481 |
| DHCPv6 guard operating mechanism | 481 |
| Restrictions and guidelines: DHCPv6 guard configuration | 482 |
| DHCPv6 guard tasks at a glance..... | 482 |
| Configuring a DHCPv6 guard policy | 482 |
| Applying a DHCPv6 guard policy to an interface | 483 |
| Applying a DHCPv6 guard policy to a VLAN..... | 483 |
| Display and maintenance commands for DHCPv6 guard..... | 483 |
| DHCPv6 guard configuration examples..... | 484 |
| Example: Configuring DHCPv6 guard..... | 484 |
| DHCPv6 guard commands | 485 |
| device-role..... | 485 |

| | |
|---|------------|
| display ipv6 dhcp guard policy | 486 |
| if-match reply acl | 488 |
| if-match server acl | 489 |
| ipv6 dhcp guard apply policy | 490 |
| ipv6 dhcp guard policy | 490 |
| preference | 491 |
| trust port | 492 |
| New feature: Ignoring router IDs during optimal route selection | 493 |
| Ignoring router IDs during optimal route selection | 493 |
| Command reference | 493 |
| bestroute router-id-ignore | 493 |
| New feature: Redistributing the local network specified in the public instance or a VPN instance | 494 |
| Redistributing the local network specified in the public instance or a VPN instance | 494 |
| Command reference | 495 |
| import-route | 495 |
| New feature: Apply route update interval setting to withdrawn routes | 496 |
| Apply route update interval setting to withdrawn routes | 496 |
| Command reference | 496 |
| route-update-interval withdrawn enable | 496 |
| New feature: Enabling prioritized withdrawal of specific routes | 497 |
| Enabling prioritized withdrawal of specific routes | 497 |
| Command reference | 498 |
| update-first route-policy | 498 |
| New feature: Configuring BGP RPKI | 499 |
| Configuring BGP RPKI | 499 |
| About BGP RPKI | 499 |
| Configuring RPKI connection parameters | 499 |
| Enabling BGP RPKI validation | 500 |
| Applying the BGP RPKI validation state to optimal route selection | 501 |
| Advertising BGP RPKI validation state to a peer or peer group | 502 |
| Resetting BGP RPKI sessions | 503 |
| Display and maintenance commands for BGP security features | 503 |
| Command reference | 504 |
| rpk | 504 |

| | |
|--|------------|
| server tcp | 504 |
| port | 505 |
| passwords | 506 |
| refresh-time | 506 |
| response-time | 507 |
| purge-time | 508 |
| check-origin-validation | 508 |
| bestroute origin-as-validation | 509 |
| peer advertise origin-as-validation | 510 |
| reset bgp rpki server | 511 |
| display bgp rpki server | 512 |
| display bgp rpki table | 514 |
| New feature: Configuring BGP policy accounting..... | 515 |
| Configuring BGP policy accounting | 515 |
| Command reference | 516 |
| bgp-policy accounting | 516 |
| display bgp-policy statistics | 517 |
| New feature: Configuring BGP to immediately send updates for routes that match an IPv6 prefix list..... | 519 |
| Configuring BGP to immediately send updates for routes that match an IPv6 prefix list..... | 519 |
| Command reference | 519 |
| bgp update-delay on-startup ipv6-prefix-list..... | 519 |
| New feature: BGP optimal route selection delay | 520 |
| Configuring BGP optimal route selection delay | 520 |
| Command reference | 521 |
| route-select delay..... | 521 |
| New feature: Specifying a VPN instance for the BMP server | 523 |
| Specifying a VPN instance for the BMP server..... | 523 |
| Command reference | 523 |
| server vpn-instance..... | 523 |
| New feature: Routing policy change delay | 524 |
| Configuring routing policy change delay | 524 |
| Command reference | 524 |
| route-policy-change delay-time | 524 |

| | |
|---|-----|
| New feature: Support of a routing policy for L3VNI configuration | 525 |
| Specifying an L3VNI in if-match and apply clauses for a routing policy..... | 525 |
| Command reference | 525 |
| apply l3-vni | 525 |
| if-match l3-vni..... | 526 |
| New feature: Configuring the BGP RPKI validation state match criterion for a routing policy node..... | 527 |
| Configuring the BGP RPKI validation state match criterion for a routing policy node..... | 527 |
| Command reference | 527 |
| if-match rpki..... | 527 |
| New feature: Enabling logging for source MAC-based ARP attack detection | 528 |
| Enabling logging for source MAC-based ARP attack detection | 528 |
| Command reference | 528 |
| arp source-mac log enable..... | 528 |
| New feature: Enabling ND attack detection logging | 529 |
| Enabling ND attack detection logging | 529 |
| Command reference | 529 |
| ipv6 nd detection log enable | 529 |
| New feature: Configuring BFD session flapping suppression..... | 530 |
| Configuring BFD session flapping suppression | 530 |
| Command reference | 531 |
| bfd dampening | 531 |
| New feature: Configuring the timer that delays reporting the first BFD session establishment failure to the data link layer | 532 |
| Configuring the timer that delays reporting the first BFD session establishment failure to the data link layer..... | 532 |
| About configuring the delay timer..... | 532 |
| Command reference | 532 |
| bfd detect-interface first-fail-timer..... | 532 |

| | |
|--|------------|
| New feature: Enabling special processing for BFD sessions | 533 |
| Enabling special processing for BFD sessions | 533 |
| About enabling special processing for BFD sessions | 533 |
| Command reference | 534 |
| bfd detect-interface special-processing..... | 534 |
| New feature: SBFD..... | 535 |
| Configuring SBFD | 535 |
| About SBFD | 535 |
| Restrictions and guidelines: SBFD configuration | 536 |
| Configuring the initiator | 536 |
| Configuring the reflector..... | 537 |
| Command reference | 537 |
| display sbfd session | 537 |
| sbfd local-discriminator | 542 |
| New feature: Associating Track with a control-mode BFD session | 543 |
| Associating Track with a control-mode BFD session | 543 |
| Command reference | 544 |
| track bfd ctrl..... | 544 |
| New feature: Enabling the device to send IGMP general queries upon a path change..... | 545 |
| Enabling the device to send IGMP general queries upon a path change | 545 |
| Command reference | 546 |
| send-query enable | 546 |
| New feature: Enabling the device to send MLD general queries upon a path change..... | 546 |
| Enabling the device to send MLD general queries upon a path change..... | 546 |
| Command reference | 547 |
| send-query enable | 547 |
| New feature: Enabling embedded RP | 547 |
| Command reference | 548 |

| | |
|--|-----|
| New feature: Specifying a local network to be advertised in the public instance or a VPN instance..... | 549 |
| Specifying a local network to be advertised in the public instance or a VPN instance | 549 |
| Command reference | 549 |
| network..... | 549 |
| New feature: Changing the next hop address of VPNv4 routes to a VPN address..... | 550 |
| Command reference | 551 |
| peer next-hop-vpn | 551 |
| New feature: Displaying tunnel policy information..... | 552 |
| Displaying tunnel policy information..... | 552 |
| Command reference | 552 |
| display tunnel-policy..... | 552 |
| New feature: SBFD for MPLS | 553 |
| Configuring SBFD for MPLS | 553 |
| Command reference | 554 |
| display mpls sbfd..... | 554 |
| mpls sbfd (for LSP) | 556 |
| mpls sbfd (for TE tunnel)..... | 557 |
| New feature: MPLS SR..... | 558 |
| About MPLS SR | 558 |
| MPLS SR characteristics | 558 |
| Basic concepts..... | 559 |
| How MPLS SR works..... | 559 |
| Dynamic SID allocation | 559 |
| Label forwarding entry installation based on SIDs..... | 560 |
| SRLSP setup..... | 561 |
| Packet forwarding in MPLS SR..... | 561 |
| MPLS SR and LDP interworking..... | 564 |
| Protocols and standards | 565 |
| MPLS SR tasks at a glance | 565 |
| IP traffic forwarding over SRLSPs tasks at a glance..... | 565 |
| Configuring IGP-based SID advertisement | 565 |
| IGP-based SID advertisement tasks at a glance | 565 |

| | |
|---|------------|
| Prerequisites | 565 |
| Configuring the IGP to support MPLS SR | 565 |
| Configuring prefix SIDs | 566 |
| Enabling MPLS SR adjacency label allocation for the IGP | 566 |
| Configuring the MPLS SRGB | 567 |
| Configuring the device to prefer SRLSPs in traffic forwarding | 567 |
| Display and maintenance commands for MPLS SR | 567 |
| MPLS SR configuration examples | 568 |
| Example: Configuring MPLS SR based on ISIS-advertised SIDs | 568 |
| Example: Configuring MPLS SR and LDP interworking | 575 |
| Command reference | 579 |
| display isis segment-routing adjacency | 579 |
| display isis segment-routing global-block | 580 |
| isis prefix-sid | 581 |
| segment-routing adjacency enable | 582 |
| segment-routing global-block | 583 |
| segment-routing mpls | 583 |
| segment-routing sr-prefer | 584 |
| New feature: Displaying buffer usage statistics for interfaces | 585 |
| Displaying buffer usage statistics for interfaces | 585 |
| Command reference | 585 |
| display buffer usage interface | 585 |
| New feature: Configuring an EAP profile | 587 |
| Configuring an EAP profile | 587 |
| Command reference | 587 |
| eap-profile | 587 |
| ca-file | 588 |
| method | 589 |
| New feature: AAA test | 590 |
| Configuring the AAA test feature | 590 |
| Command reference | 592 |
| exclude | 592 |
| include | 593 |
| radius attribute-test-group | 595 |
| test-aaa | 596 |

| | |
|---|-----|
| New feature: Displaying and clearing history authentication and accounting load statistics for all RADIUS servers | 599 |
| Displaying and clearing history authentication and accounting load statistics for all RADIUS servers | 599 |
| Command reference | 600 |
| display radius server-load statistics | 600 |
| reset radius server-load statistics | 601 |
| New feature: 802.1X offline detection | 602 |
| Configuring 802.1X offline detection | 602 |
| Command reference | 602 |
| dot1x offline-detect enable | 602 |
| dot1x timer offline-detect | 603 |
| New feature: 802.1X unauthenticated user aging | 604 |
| Configuring 802.1X unauthenticated user aging | 604 |
| Command reference | 605 |
| dot1x unauthenticated-user aging enable | 605 |
| dot1x timer user-aging | 606 |
| New feature: 802.1X EAP-TLS fragmentation for packets sent to the server | 607 |
| Setting the maximum length of an EAP-TLS fragment sent to the server | 607 |
| Command reference | 607 |
| dot1x eap-tls-fragment to-server | 607 |
| New feature: User aging for unauthenticated MAC authentication users ... | 608 |
| Configuring user aging for unauthenticated MAC authentication users | 608 |
| Command reference | 609 |
| mac-authentication unauthenticated-user aging enable | 609 |
| mac-authentication timer user-aging | 610 |
| New feature: MAC authentication method | 611 |
| Specifying a MAC authentication method | 611 |
| Command reference | 611 |
| mac-authentication authentication-method | 611 |

| | |
|--|------------|
| New feature: User-specific MAC authentication offline detection | 612 |
| Configuring offline detection for a specific MAC authentication user | 612 |
| Command reference | 613 |
| mac-authentication offline-detect mac-address | 613 |
| New feature: NAS-Port-Type attribute | 614 |
| Configuring the NAS-Port-Type attribute | 614 |
| Command reference | 615 |
| portal nas-port-type | 615 |
| New feature: Checking the issuing of category-2 portal filtering rules | 616 |
| Checking the issuing of category-2 portal filtering rules..... | 616 |
| Command reference | 616 |
| portal user-rule assign-check enable | 616 |
| New feature: Configuring user authentication timeout timer for global password control..... | 617 |
| Configuring user authentication timeout timer for global password control..... | 617 |
| About configuring user authentication timeout timer for global password control | 617 |
| Command reference | 618 |
| password-control authentication-timeout | 618 |
| New feature: Configuring password control for local network access users | 618 |
| Setting global password control features for network access users | 618 |
| Setting password control parameters for a single network access user | 620 |
| Command reference | 620 |
| Modified command: display password-control | 620 |
| Modified feature: password-control complexity | 622 |
| Modified feature: password-control composition..... | 622 |
| Modified feature: password-control enable | 623 |
| Modified feature: password-control history | 623 |
| Modified feature: password-control length | 623 |
| Modified feature: password-control update-interval | 624 |
| Modified feature: reset password-control history-record | 624 |
| New feature: Configuring password change at first login | 624 |
| Configuring password change at first login | 624 |
| Command reference | 625 |

| | |
|---|------------|
| password-control change-password first-login enable | 625 |
| New feature: Enabling the SSL server to send the complete certificate chain during SSL negotiation | 626 |
| Enabling the SSL server to send the complete certificate chain during SSL negotiation..... | 626 |
| Command reference | 626 |
| certificate-chain-sending enable | 626 |
| New feature: Support of RRPP for the link-up delay timer | 627 |
| Configuring the link-up delay timer..... | 627 |
| Command reference | 628 |
| linkup-delay-timer | 628 |
| New feature: gRPC..... | 628 |
| About gRPC | 628 |
| gRPC protocol stack layers..... | 629 |
| Network architecture | 629 |
| Telemetry technology based on gRPC | 629 |
| Telemetry modes | 629 |
| Protocols | 630 |
| FIPS compliance..... | 630 |
| Configuring the gRPC dial-in mode..... | 630 |
| gRPC dial-in mode configuration tasks at a glance | 630 |
| Configuring the gRPC service..... | 630 |
| Configuring a gRPC user | 630 |
| Configuring the gRPC dial-out mode | 631 |
| gRPC dial-out mode configuration tasks at a glance | 631 |
| Enabling the gRPC service | 631 |
| Configuring sensors | 631 |
| Configuring collectors..... | 632 |
| Configuring a subscription..... | 632 |
| Display and maintenance commands for gRPC | 633 |
| gRPC configuration examples | 633 |
| Example: Configuring the gRPC dial-in mode..... | 633 |
| Example: Configuring the gRPC dial-out mode | 634 |
| gRPC dial-in mode commands | 635 |
| display grpc | 635 |
| grpc enable | 636 |
| grpc idle-timeout..... | 636 |
| grpc port | 637 |
| gRPC dial-out mode commands | 638 |

| | |
|--|-----|
| destination-group (subscription view)..... | 638 |
| destination-group (telemetry view)..... | 638 |
| ipv4-address..... | 639 |
| ipv6-address..... | 640 |
| sensor path | 641 |
| sensor-group (subscription view) | 641 |
| sensor-group (telemetry view)..... | 642 |
| source-address | 643 |
| subscription | 643 |
| telemetry | 644 |

New feature: Protocol buffers 645

| | |
|---|-----|
| Protocol buffers code format..... | 645 |
| Proto definition files..... | 646 |
| Proto definition files in dial-in mode | 646 |
| Proto definition file in dial-out mode..... | 647 |
| Obtaining proto definition files..... | 648 |
| Example: Developing a gRPC collector-side application | 648 |
| Prerequisites | 648 |
| Generating the C++ code for the proto definition files..... | 648 |
| Developing the collector-side application..... | 649 |

New feature: Specifying the NTP time-offset thresholds for log and trap outputs 653

| | |
|---|-----|
| Specifying the NTP time-offset thresholds for log and trap outputs..... | 653 |
| Command reference | 654 |
| ntp-service time-offset-threshold..... | 654 |

New feature: Specifying the SNTP time-offset thresholds for log and trap outputs..... 654

| | |
|--|-----|
| Specifying the SNTP time-offset thresholds for log and trap outputs..... | 654 |
| Command reference | 655 |
| sntp time-offset-threshold..... | 655 |

New feature: Ignoring the permit flag added by OpenFlow 656

| | |
|--|-----|
| Ignoring the permit flag added by OpenFlow | 656 |
| About ignoring the permit flag added by OpenFlow | 656 |
| Procedure..... | 656 |
| Command reference | 656 |

| | |
|--|------------|
| openflow permit-flag ignore..... | 656 |
| New feature: ND flood suppression | 657 |
| Enabling ND flood suppression..... | 657 |
| Command reference | 657 |
| display ipv6 nd suppression vsi..... | 657 |
| ipv6 nd suppression enable | 658 |
| reset ipv6 nd suppression vsi..... | 659 |
| New feature: IP address configuration of the VTEPs in a DR system | 659 |
| Specifying the IP addresses of the VTEPs in a DR system | 659 |
| Command reference | 660 |
| evpn drni local | 660 |
| New feature: Redistribution of MAC/IP advertisement routes into BGP unicast routing tables | 661 |
| Redistributing MAC/IP advertisement routes into BGP unicast routing tables..... | 661 |
| Command reference | 662 |
| import evpn mac-ip..... | 662 |
| New feature: Route exchange between the BGP EVPN address family and the BGP VPNv4 or VPNv6 address family | 663 |
| Configuring the BGP EVPN address family and the BGP VPNv4 or VPNv6 address family to exchange routes | 663 |
| About route exchange | 663 |
| Enabling BGP VPNv4 or VPNv6 route advertisement for the BGP EVPN address family | 663 |
| Enabling BGP EVPN route advertisement for the BGP VPNv4 or VPNv6 address family | 664 |
| Command reference | 664 |
| advertise l3vpn route..... | 664 |
| advertise evpn route..... | 665 |
| New feature: Disabling the ACLs issued by the OVSDB controller | 666 |
| Disabling the ACLs issued by the OVSDB controller | 666 |
| Command reference | 666 |
| vtep acl disable | 666 |

| | |
|--|-----|
| New feature: Replacement of the L3 VXLAN ID and RD of IP prefix advertisement routes on EDs..... | 667 |
| Enabling an ED to replace the L3 VXLAN ID and RD of IP prefix advertisement routes..... | 667 |
| Command reference | 668 |
| peer re-originated..... | 668 |
| New feature: BGP EVPN route advertisement suppression | 669 |
| Suppressing BGP EVPN route advertisement | 669 |
| Command reference | 669 |
| peer advertise evpn-route suppress..... | 669 |
| New feature: Removal of the default-gateway extended community attribute from EVPN gateway routes..... | 670 |
| Removing the default-gateway extended community attribute from EVPN gateway routes | 670 |
| Command reference | 671 |
| peer default-gateway no-advertise..... | 671 |
| New feature: Configuring a 16/32-Gbps FC interface | 672 |
| Configuring a 16/32-Gbps FC interface | 672 |
| About configuring a 16/32-Gbps FC interface..... | 672 |
| Configuring the speed of a 16/32-Gbps FC interface..... | 672 |
| Splitting a 32-Gbps FC interface and combining 16-Gbps FC breakout interfaces | 672 |
| Changing the port type between a Layer 2 Ethernet interface and an FC interface | 673 |
| Command reference | 674 |
| New command: using..... | 674 |
| Modified command: speed | 674 |
| New feature: Enabling area ID-to-F_Port binding..... | 675 |
| Enabling area ID-to-F_Port binding..... | 675 |
| fc port-bind-area enable | 675 |
| display fc port-bind-area..... | 676 |
| New feature: BB_Credit recovery for FC interfaces | 677 |
| Enabling BB_Credit recovery for an FC interface | 677 |
| Command reference | 678 |
| fcb2bcredit recovery enable | 678 |

| | |
|--|-----|
| Modified feature: Software upgrade | 678 |
| Feature change description..... | 678 |
| Command changes | 679 |
| Modified feature: Using issu commands for an ISSU | 679 |
| Feature change description..... | 679 |
| Command changes | 679 |
| Modified command: display version comp-matrix | 679 |
| Modified command: issu load..... | 679 |
| Modified command: issu one-step | 680 |
| Modified feature: Per-packet load sharing mode of aggregation groups | 681 |
| Feature change description..... | 681 |
| Command changes | 681 |
| Modified command: link-aggregation load-sharing mode | 681 |
| Modified feature: Configuring MAC-to-VLAN entries | 681 |
| Feature change description..... | 681 |
| Command changes | 681 |
| Modified command: mac-vlan mac-address | 681 |
| Modified feature: Configuring the advertisable TLVs..... | 682 |
| Feature change description..... | 682 |
| Command changes | 682 |
| Modified command: lldp tlv-enable..... | 682 |
| Modified feature: Specifying the management address advertised in global management address TLV advertisement setting | 685 |
| Feature change description..... | 685 |
| Command changes | 686 |
| Modified command: lldp global tlv-enable basic-tlv management-address-tlv | 686 |
| Modified feature: Configuring aging timer for dynamic ARP entries | 686 |
| Feature change description..... | 686 |
| Command changes | 686 |
| Modified command: arp timer aging..... | 686 |

| | |
|---|-----|
| Modified feature: Enabling ARP snooping..... | 687 |
| Feature change description..... | 687 |
| Command changes..... | 687 |
| Modified command: arp snooping enable | 687 |
| Modified feature: Displaying ARP snooping entries | 688 |
| Feature change description..... | 688 |
| Command changes..... | 688 |
| Modified command: display arp snooping..... | 688 |
| Modified feature: Clearing ARP snooping entries..... | 689 |
| Feature change description..... | 689 |
| Command changes..... | 689 |
| Modified command: reset arp snooping | 689 |
| Modified feature: Configuring DHCP snooping handling strategy for Option 82 in request messages..... | 689 |
| Feature change description..... | 689 |
| Command changes | 690 |
| Modified command: dhcp snooping information strategy..... | 690 |
| Modified feature: Setting the interface MTU for IPv6 packets | 690 |
| Feature change description..... | 690 |
| Command changes | 690 |
| Modified command: ipv6 mtu | 690 |
| Modified feature: Support for specifying next hop indexes in static routes · | 691 |
| Feature change description..... | 691 |
| Command changes..... | 691 |
| Modified command: ip route-static | 691 |
| Modified feature: Support for configuring OSPF commands in VSI interface view | 692 |
| Feature change description..... | 692 |
| Modified feature: Displaying IS-IS LSP log information | 693 |
| Feature change description..... | 693 |

| | |
|---|------------|
| Command changes | 693 |
| Modified command: display isis event-log lsp | 693 |
| Modified feature: Clearing IS-IS LSP log information | 693 |
| Feature change description | 693 |
| Command changes | 693 |
| Modified command: reset isis event-log lsp | 693 |
| Modified feature: Filtering advertised/received BGP routes | 694 |
| Feature change description | 694 |
| Command changes | 694 |
| Modified command: filter-policy export/filter-policy import | 694 |
| Modified feature: Filtering routes advertised to or received from a peer or peer group by using an ACL | 695 |
| Feature change description | 695 |
| Command changes | 695 |
| Modified command: peer filter-policy | 695 |
| Modified feature: Configuring BMP | 697 |
| Feature change description | 697 |
| Command changes | 697 |
| New command: server connect-interface | 697 |
| New command: route-mode adj-rib-out | 698 |
| New command: route-mode loc-rib | 698 |
| Modified feature: Displaying detailed BGP routing information | 699 |
| Feature change description | 699 |
| Command changes | 699 |
| Modified command: display bgp routing-table ipv4 multicast | 699 |
| Modified command: display bgp routing-table ipv4 rfilter | 701 |
| Modified command: display bgp routing-table ipv4 unicast | 703 |
| Modified command: display bgp routing-table ipv6 multicast | 706 |
| Modified command: display bgp routing-table ipv6 unicast | 707 |
| Modified feature: Applying a routing policy to routes outgoing to a peer or peer group | 709 |
| Feature change description | 709 |

| | |
|--|------------|
| Command changes | 710 |
| Modified feature: Specifying an ACL | 710 |
| Feature change description | 710 |
| Command changes | 710 |
| Modified feature: Defining an ACL match criterion in a traffic class of a QoS | |
| policy | 710 |
| Feature change description | 710 |
| Command changes | 710 |
| Modified command: if-match | 710 |
| Modified feature: Applying a QoS policy globally | 711 |
| Feature change description | 711 |
| Command changes | 711 |
| Modified command: qos apply policy global | 711 |
| Modified command: display qos policy global | 711 |
| Modified command: display qos policy diagnosis global | 712 |
| Modified command: reset qos policy global | 712 |
| Modified feature: Configuring a test profile for RADIUS server status detection | |
| | 712 |
| Feature change description | 712 |
| Command changes | 713 |
| Modified command: radius-server test-profile | 713 |
| Modified feature: RADIUS server quiet timer | 713 |
| Feature change description | 713 |
| Command changes | 713 |
| Modified command: timer quiet (RADIUS scheme view) | 713 |
| Modified feature: Specifying the source IP address for outgoing RADIUS | |
| packets | 714 |
| Feature change description | 714 |
| Command changes | 714 |
| Modified command: nas-ip (RADIUS scheme view) | 714 |
| Modified command: radius nas-ip | 715 |

| | |
|---|-----|
| Modified feature: Specifying the source IP address for outgoing HWTACACS packets | 715 |
| Feature change description..... | 715 |
| Command changes | 715 |
| Modified command: nas-ip (HWTACACS scheme view) | 715 |
| Modified command: hwtacacs nas-ip | 716 |
| Modified feature: Including user IP addresses in MAC authentication requests | 717 |
| Feature change description..... | 717 |
| Command changes | 717 |
| Modified command: mac-authentication carry user-ip | 717 |
| Modified feature: Configuring MAC-based MAC authentication user accounts | 718 |
| Feature change description..... | 718 |
| Command changes | 718 |
| Modified command: mac-authentication user-name-format..... | 718 |
| Modified feature: Handling new MAC access attempts in a VLAN after port security's MAC address limit for that VLAN is reached | 718 |
| Feature change description..... | 718 |
| Command changes | 719 |
| Modified feature: Port security NTK feature | 719 |
| Feature change description..... | 719 |
| Command changes | 719 |
| Modified command: port-security ntk-mode | 719 |
| Modified feature: Password handling manners with password control enabled globally | 720 |
| Feature change description..... | 720 |
| Password handling manners for device management users | 720 |
| Password handling manners for obtaining temporary user role authorization | 721 |

| | |
|--|------------|
| Command changes | 721 |
| Modified feature: RSA key modulus length | 721 |
| Feature change description | 721 |
| Command changes | 721 |
| Modified command: public-key local create | 721 |
| Modified feature: Key modulus length of the RSA key pair used for certificate request in a PKI domain | 722 |
| Feature change description | 722 |
| Command changes | 722 |
| Modified command: public-key rsa | 722 |
| Modified feature: Configuring the ECDSA signature authentication method in an IKE proposal | 723 |
| Feature change description | 723 |
| Command changes | 723 |
| Modified command: authentication-method | 723 |
| Modified feature: Specifying the DH group used for key negotiation in IKE phase 1 | 723 |
| Feature change description | 723 |
| Command changes | 724 |
| Modified command: dh | 724 |
| Modified feature: Keyboard-interactive authentication support for SSH users | 724 |
| Feature change description | 724 |
| Command changes | 725 |
| Modified command: ssh user | 725 |
| Modified feature: Displaying IPv4 source guard bindings | 725 |
| Feature change description | 725 |
| Command changes | 726 |

| | |
|--|-----|
| Modified feature: Displaying IPv6 source guard bindings | 726 |
| Feature change description..... | 726 |
| Command changes | 726 |
| Modified feature: Configuring ARP attack detection logging | 727 |
| Feature change description..... | 727 |
| Command changes | 727 |
| Modified command: arp detection log enable | 727 |
| Modified feature: Creating a BFD session for detecting the local interface state | |
| | 728 |
| Feature change description..... | 728 |
| Command changes | 728 |
| Modified command: bfd detect-interface | 728 |
| Modified feature: Configuring the BFD authentication mode for single-hop BFD | |
| control packets | 728 |
| Feature change description..... | 728 |
| Command changes | 729 |
| Modified command: bfd authentication-mode | 729 |
| Modified feature: Setting the minimum interval for receiving BFD echo packets | |
| | 729 |
| Feature change description..... | 729 |
| Command changes | 729 |
| Modified command: bfd min-echo-receive-interval | 729 |
| Modified feature: Support for specific BFD commands in VSI interface view | |
| | 730 |
| Feature change description..... | 730 |
| Command changes | 730 |

Modified feature: Setting the maximum number of active routes in a VPN

instance 730

Feature change description..... 730

Command changes 730

Modified command: routing-table limit 730

Modified feature: Enabling BFD on an MPLS TE tunnel interface 731

Feature change description..... 731

Command changes 731

Modified command: mpls bfd 731

Modified feature: Setting the length of ICMP or ICMPv6 echo requests for the

ping operation 731

Feature change description..... 731

Command changes 731

Modified command: ping 731

Modified command: ping ipv6 732

Modified feature: Specifying a source interface for NTP messages 732

Feature change description..... 732

Command changes 732

Modified command: ntp-service source 732

Modified feature: PTP 733

Feature change description..... 733

Command changes 733

Modified command: ptp profile 733

Modified command: display ptp clock 733

Modified command: ptp announce-interval 735

Modified command: ptp clock-source..... 736

Modified command: ptp domain 737

Modified command: ptp min-delayreq-interval 738

Modified command: ptp pdelay-req-interval 738

Modified command: ptp syn-interval 739

Modified feature: Shutting down an interface by using OpenFlow 739

Feature change description..... 739

| | |
|--|------------|
| Command changes | 739 |
| Modified command: openflow shutdown | 739 |
| Modified feature: Configuring flow sampling of sFlow | 740 |
| Feature change description | 740 |
| Command changes | 740 |
| Modified command: sflow counter collector | 740 |
| Modified feature: Configuring counter sampling of sFlow | 740 |
| Feature change description | 740 |
| Command changes | 741 |
| Modified command: sflow counter collector | 741 |
| Modified feature: Testing the reachability of remote VMs in VXLANs | 741 |
| Feature change description | 741 |
| Command changes | 741 |
| New command: emulate-ping vxlan enable | 741 |
| Modified feature: Enabling packet statistics for automatically created VXLAN tunnels | 742 |
| Feature change description | 742 |
| Command changes | 742 |
| Modified command: tunnel statistics vxlan auto | 742 |
| Modified feature: Disabling flooding for a VSI | 743 |
| Feature change description | 743 |
| Command changes | 743 |
| Modified command: flooding disable | 743 |
| Modified feature: Setting the maximum bandwidth for an AC | 743 |
| Feature change description | 743 |
| Command changes | 744 |
| Modified command: bandwidth | 744 |
| Modified feature: Displaying information about peers automatically discovered through BGP in EVPN | 744 |
| Feature change description | 744 |
| Command changes | 744 |

| | |
|---|------------|
| Modified command: display evpn auto-discovery | 744 |
| Modified feature: PW redundancy..... | 745 |
| Feature change description..... | 745 |
| Command changes | 745 |
| New command: pw-redundancy..... | 745 |
| Modified command: display l2vpn ldp | 746 |
| Modified command: display l2vpn xconnect-group | 747 |
| Modified command: display l2vpn vsi..... | 747 |
| Modified command: peer..... | 748 |
| Modified feature: Displaying log buffer information and buffered logs | 749 |
| Feature change description..... | 749 |
| Command changes | 749 |
| Modified command: display logbuffer..... | 749 |
| Modified feature: Configuring the timestamp for logs sent to log hosts | 750 |
| Feature change description..... | 750 |
| Command changes | 750 |
| Modified command: display logbuffer..... | 750 |
| Modified feature: Configuring log output to a log host | 750 |
| Feature change description..... | 750 |
| Command changes | 750 |
| Modified command: info-center loghost | 750 |
| Modified feature: IRF master election during automated VCF fabric deployment | |
| | 751 |
| Feature change description..... | 751 |
| Command changes | 751 |

About software feature changes

This document contains feature changes in the software versions listed in [Table 1](#). For information about software feature changes in HPE 5940-CMW710-R2612P02 (or earlier), see their respective release notes (software feature changes).

Table 1 Software feature change summary

| Section | Software feature changes |
|-----------------|---|
| Release 6710P03 | Contains changes in HPE 5940-CMW710-R6710P03 over HPE 5940-CMW710-R6710P01. |
| Release 6710P01 | Contains changes in HPE 5940-CMW710-R6710P01 over HPE 5940-CMW710-R6710. |
| Release 6710 | Contains changes in HPE 5940-CMW710-R6710 over HPE 5940-CMW710-E6705. |
| ESS 6705 | Contains changes in HPE 5940-CMW710-E6705 over HPE 5940-CMW710-E6702. |
| ESS 6702 | Contains changes in HPE 5940-CMW710-E6702 over HPE 5940-CMW710-R2717. |
| Release 2717 | Contains changes in HPE 5940-CMW710-R2717 over HPE 5940-CMW710-R2702. |
| Release 2702 | Contains changes in HPE 5940-CMW710-R2702 over HPE 5940-CMW710-F2701. |
| Feature 2701 | Contains changes in HPE 5940-CMW710-F2701 over HPE 5940-CMW710-R2612P02. |

Version R6710P03

This release has the following changes:

- New feature: DRNI configuraton
- New feature: Generating a log message to display product version numbers before and after a software upgrade
- Modified feature: Enabling SNMP notification for FC interface status changes
- Modified feature: Adding a member to a zone alias
- Modified feature: Adding a member to a zone
- Modified feature: Displaying parent information for a zone member

New feature: DRNI configuraton

All DRNI commands were newly added.

New feature: Generating a log message to display product version numbers before and after a software upgrade

Generating a log message to display product version numbers before and after a software upgrade

As from this version, the device can generate a log message to display the product version numbers before and after a software upgrade when it successfully upgrades the software by executing the `install`, `issu`, or `boot-loader` command.

Log message reference

PKG_UPGRADE_INFO

| | |
|---------------------------|--|
| Message text | The [STRING] device upgraded the software version from [STRING] to software [STRING]. |
| Variable fields | \$1: Device name. \$2: Version number before software upgrade. \$3: Version number after software upgrade. |
| Severity level | 5 (Notification) |
| Example | PKG/5/PKG_UPGRADE_INFO: The HPE FF 5940 32QSFP+ Switch device upgraded the software version from software version 1-patch version 1 to software version 2-patch version 2. |
| Impact | No negative impact on the system. |
| Cause | Executed the install , issu , or boot-loader command successfully and the new version took effect. |
| Recommended action | No action is required. |

Modified feature: Enabling SNMP notification for FC interface status changes

Feature change description

As from this release, the system supports enabling SNMP notification for status changes of an FC interface.

Command changes

Modified command: enable snmp trap updown

```
enable snmp trap updown
```

Views

Interface view

Change description

Before modification: The system does not support enabling SNMP notification for status changes of an FC interface.

After modification: The system supports enabling SNMP notification for status changes of an FC interface.

Modified feature: Adding a member to a zone alias

Feature change description

As from this release, the system supports adding an interface as a member of a zone alias.

Command changes

Modified command: `member` (zone alias view)

Old syntax

```
member { fcid fcid | fwwn fwwn | pwwn pwwn } [ initiator | target ]  
undo member { fcid fcid | fwwn fwwn | pwwn pwwn }
```

New syntax

```
member { fcid fcid | fwwn fwwn | interface interface-type interface-number  
| pwwn pwwn } [ initiator | target ]  
undo member { fcid fcid | fwwn fwwn | interface interface-type  
interface-number | pwwn pwwn }
```

Views

Zone alias view

Change description

Before modification: You cannot add an interface as a member of a zone alias.

After modification: You can add an interface as a member of a zone alias.

interface *interface-type interface-number*: Specifies an interface by its type and number, which can only be an FC interface.

Modified feature: Adding a member to a zone

Feature change description

As from this release, the system supports adding an interface as a member of a zone.

Command changes

Modified command: `member` (zone view)

Old syntax

```
member { { fcid fcid | fwwn fwwn | pwwn pwwn } [ initiator | target ] |  
zone-alias zone-alias-name }  
undo member { fcid fcid | fwwn fwwn | pwwn pwwn | zone-alias  
zone-alias-name }
```

New syntax

```
member { { fcid fcid | fwwn fwwn | interface interface-type
interface-number | pwwn pwwn } [ initiator | target ] | zone-alias
zone-alias-name }

undo member { fcid fcid | fwwn fwwn | interface interface-type
interface-number | pwwn pwwn | zone-alias zone-alias-name }
```

Views

Zone view

Change description

Before modification: You cannot add an interface as a member of a zone.

After modification: You can add an interface as a member of a zone.

interface *interface-type interface-number*: Specifies an interface by its type and number, which can only be an FC interface.

Modified feature: Displaying parent information for a zone member

Feature change description

As from this release, the system supports displaying parent information for an interface member.

Command changes

Modified command: display zone member

Old syntax

```
display zone member { fcid fcid | fwwn fwwn | pwwn pwwn | zone-alias
zone-alias-name } [ vsan vsan-id ]
```

New syntax

```
display zone member { fcid fcid | fwwn fwwn | interface interface-type
interface-number | pwwn pwwn | zone-alias zone-alias-name } [ vsan
vsan-id ]
```

Views

Zone view

Change description

Before modification: You cannot display parent information for an interface member.

After modification: You can display parent information for an interface member.

interface *interface-type interface-number*: Specifies an interface by its type and number, which can only be an FC interface.

Release 6710P01

This release has no feature changes.

Release 6710

This release has the following changes:

- New features: Fundamentals features
- New features: Layer 2—LAN switching features
- New features: Layer 3—IP services features
- New features: Layer 3—IP routing features
- New features: IP multicast features
- New features: MPLS features
- New features: ACL and QoS features
- New features: Security features
- New features: High availability features
- New features: Network management and monitoring features
- New features: Telemetry features
- New features: OpenFlow features
- New features: VXLAN features
- New features: EVPN features
- New features: Intelligent lossless network features
- New features: M-LAG support for DRNI commands
- New feature: EVPN instance created in system view
- Modified feature: IPv6 routes with prefixes longer than 64 bits
- Modified feature: Match criteria in a traffic class
- Modified feature: Associating a traffic behavior with a traffic class
- Modified feature: Displaying the running configuration
- Modified feature: Displaying the contents of the configuration file for the next system startup
- Modified feature: Optimized display of BGP BMP server information
- Modified feature: Disabling BGP session establishment with peers and peer groups
- Modified feature: Optimizations to VXLAN command output
- Modified feature: Sharing of VSI interfaces among VSIs
- Modified feature: Enabling L2TP for the specified protocol
- Modified feature: Creating a service loopback group and specify its service type
- Modified feature: Displaying LSP information
- Modified feature: Configuring the signaling protocol for an MPLS TE tunnel interface
- Modified feature: Configuring the tunnel selection order and set the number of tunnels for load balancing
- Modified feature: Verifying MPLS LSP connectivity for an IPv4 prefix
- Modified feature: Tracing MPLS LSPs from the ingress node to the egress node for an IPv4 prefix
- Modified feature: Creating a local site
- Modified feature: Enabling link flapping protection on an interface
- Modified feature: AAA methods in an ISP domain
- Modified feature: Setting the 802.1X periodic reauthentication timer

- Modified feature: Setting the periodic MAC reauthentication timer
- Modified feature: Enabling SNMP notifications for L3VPN
- Modified feature: Displaying BGP VPNv4 routing information
- Modified feature: Displaying BGP VPNv6 routing information
- Modified feature: Creating an SNMPv3 user
- Modified feature: Displaying local public keys
- Modified feature: Flow-mirroring traffic to an interface
- Modified feature: Disabling the device from learning the source MAC addresses of Layer 2 protocol packets
- Modified feature: Restrictions for an interface to join a Layer 2 aggregation group
- Modified feature: Specifying multiple VLAN IDs for VLAN name TLVs
- Modified feature: Applying a QoS policy to an interface

New features: Fundamentals features

Table 1 describes the fundamentals features added in this software version.

For more information about the features, see *Fundamentals Configuration Guide* in *HPE 5940 Switch Series Configuration Guides-Release 671x*.

For more information about the commands, see *Fundamentals Command Reference* in *HPE 5940 Switch Series Command References-Release 671x*.

Table 1 Fundamentals features in Release 6710

| Feature | Command changes |
|--|---|
| Login management: Controlling TCP connections from IPv6 HTTP and HTTPS clients. | The following commands were added: <ul style="list-style-type: none"> • http ipv6 acl • https ipv6 acl |
| Login management: Specifying the service port number for RESTful access over HTTP/HTTPS. | The following commands were added: <ul style="list-style-type: none"> • restful http port • restful https port |
| Login management: Applying an SSL server policy to the RESTful access over HTTPS service. | The restful https ssl-server-policy command was added. |
| FTP and TFTP: Enabling the TFTP server. | The tftp server enable command was added. |
| FTP and TFTP: Setting the TFTP server working directory. | The tftp server work-directory command was added. |
| Configuration file management: Assigning a user exclusive write access to the configuration. | The following commands were added: <ul style="list-style-type: none"> • configuration exclusive by-user-name • display configuration exclusive by-user-name |
| Device management: Displaying the historical power consumption information in a coordinate system. | The display power history command was added. |
| Device management: Displaying system health status information. | The display system health command was added. |
| Device management: Displaying historical | The display system health history command |

| Feature | Command changes |
|--|--|
| system health status change information. | was added. |
| Device management: Enabling periodic CPU usage logging. | The monitor cpu-usage logging interval command was added. |
| Device management: Enabling periodic memory usage logging. | The monitor memory-usage logging interval command was added. |
| Device management: Configuring error notifications for software and hardware table entry consistency check. | The following commands were added: <ul style="list-style-type: none"> • parity-error consistency-check log enable • parity-error consistency-check threshold |
| Device management: Configuring unrecoverable error notifications for critical hardware entry parity check and ECC check. | The following commands were added: <ul style="list-style-type: none"> • parity-error unrecoverable log enable • parity-error unrecoverable period • parity-error unrecoverable reboot • parity-error unrecoverable threshold |
| Device management: Enabling transceiver monitoring. | The following commands were added: <ul style="list-style-type: none"> • transceiver monitor enable • transceiver monitor interval |
| Device management: Displaying or saving IFMGR-related operating information. | The ifmgr keyword was added to the display diagnostic-information command. |
| GIR features | All GIR commands were newly added. |

New features: Layer 2—LAN switching features

Table 2 describes the Layer 2—LAN switching features added in this software version.

For more information about the features, see *Layer 2—LAN Switching Configuration Guide* in *HPE 5940 Switch Series Configuration Guides-Release 671x*.

For more information about the commands, see *Layer 2—LAN Switching Command Reference* in *HPE 5940 Switch Series Command References-Release 671x*.

Table 2 Layer 2—LAN switching features added in Release 6710

| Feature | Command changes |
|--|---|
| Ethernet interface: Displaying the physical link state change statistics of interfaces | The display link-state-change statistics interface command was added. |
| Ethernet interface: Clearing the physical link state change statistics of interfaces | The reset link-state-change statistics interface command was added. |
| Ethernet interface: Configuring alarm parameters for sent or received pause frames | The following commands were added: <ul style="list-style-type: none"> • ifmonitor input-usage • ifmonitor output-usage • ifmonitor rx-pause • ifmonitor tx-pause • port ifmonitor input-usage |

| Feature | Command changes |
|---|---|
| | <ul style="list-style-type: none"> <code>port ifmonitor output-usage</code> <code>port ifmonitor rx-pause</code> <code>port ifmonitor tx-pause</code> |
| Ethernet interface: Shutting down all physical interfaces except the management Ethernet interfaces | The shutdown all-physical-interfaces command was added. |
| VLAN: Configuring the 802.1p priority for control packets sent by the device | The control-packet dot1p priority command was added. |
| VLAN mapping: Configuring many-to-one VLAN mapping in the IPv6 address environment | None. |
| Loop detection: Setting the loop protection delay timer | <p>The following commands were added:</p> <ul style="list-style-type: none"> loopback-detection global delay-timer time loopback-detection delay-timer time |
| Loop detection: Configuring loop detection in a VXLAN network | <p>The following commands were added:</p> <ul style="list-style-type: none"> display loopback-detection loopback-detection action { block shutdown } loopback-detection enable [vlan vlan-id-list] loopback-detection enable s-vid vlan-id-list c-vid vlan-id-list loopback-detection interval-time interval loopback-detection priority priority |
| Spanning tree: Enabling BPDU filter | <p>The following commands were added:</p> <ul style="list-style-type: none"> stp bpdu-filter stp port bpdu-filter { disable enable } |
| Spanning tree: Enabling loopback guard on an interface | The stp loopback-protection command was added. |
| LLDP: Setting the port ID subtype of port ID TLVs advertised by LLDP on a device | The lldp tlv-config basic-tlv port-id type-id command was added. |
| L2PT: Configuring user-defined Layer 2 protocols | The l2protocol user-defined protocol-name protocol-mac protocol-mac [encap-type { { ethernetii snap } protocol-type protocol-type-value llc dsap dsap-value ssap ssap-value }] command was added. |
| Ethernet link aggregation: Displaying the aggregation states of aggregation member ports and the reason why a port was placed in Unselected state | The display link-aggregation troubleshooting command was added. |
| Ethernet link aggregation: Enabling automatic link aggregation | The link-aggregation auto-aggregation enable command was added. |

| Feature | Command changes |
|--|---|
| Ethernet link aggregation: Isolating aggregate interfaces on the device | The link-aggregation lacp isolate command was added. |
| Ethernet link aggregation: Configuring LACP system settings on an aggregate interface | The following commands were added: <ul style="list-style-type: none"> • port lacp system-mac • port lacp system-priority |
| Ethernet link aggregation: Displaying information about all member ports in an aggregation group | The all-configuration keyword was added to the display link-aggregation verbose command. |
| M-LAG: Displaying the configuration consistency check status | The display m-lag consistency-check status command was added. |
| M-LAG: Displaying M-LAG troubleshooting information | The display m-lag troubleshooting command was added. |
| M-LAG: Enabling M-LAG packet authentication and configuring an authentication key | The m-lag authentication key command was added. |
| M-LAG: Setting the mode of configuration consistency check | The m-lag consistency-check mode command was added. |
| M-LAG: Configuring M-LAG extra VLANs | The m-lag extra-vlan command was added. |
| M-LAG: Enabling the peer-link interface to retain MAC address entries for single-homed devices | The m-lag peer-link mac-address hold command was added. |
| M-LAG: Associating the keepalive link with a track entry | The m-lag keepalive track command was added. |
| M-LAG: Excluding all logical interfaces from the shutdown action by M-LAG MAD | The m-lag mad exclude logical-interfaces command was added. |
| M-LAG: Enabling M-LAG MAD DOWN state persistence | The m-lag mad persistent command was added. |
| M-LAG: Bringing up the interfaces in M-LAG MAD DOWN state | The m-lag mad restore command was added. |
| M-LAG: Enabling M-LAG sequence number check | The m-lag sequence enable command was added. |
| M-LAG: Enabling M-LAG standalone mode | The m-lag standalone enable command was added. |
| M-LAG: Clearing M-LAG troubleshooting records | The reset m-lag troubleshooting history command was added. |
| M-LAG: Configuring dynamic routing access to M-LAG | The following commands were added: <ul style="list-style-type: none"> • port m-lag virtual-ip • port m-lag ipv6 virtual-ip • display m-lag virtual-ip |
| M-LAG: Configuring M-LAG system settings on an aggregate interface | The following commands were added: <ul style="list-style-type: none"> • port m-lag system-priority • port m-lag system-mac |

New features: Layer 3—IP services features

Table 3 describes the Layer 3 IP services features added in this software version.

- For more information about the features, see *Layer 3—IP services Configuration Guide* in *HPE 5940 Switch Series Configuration Guides-Release 671x*.
- For more information about the commands, see *Layer 3—IP services Command Reference* in *HPE 5940 Switch Series Command References-Release 671x*.

Table 3 Layer 3 IP services features added in Release 6710

| Feature | Command changes |
|--|---|
| ARP: Enabling error logging for ARP entry deployment to hardware | The arp hardware log enable command was added. |
| ARP: Displaying the ARP table usage | The display arp usage command was added. |
| ARP: Displaying statistics about proxy ARP reply packets | The display proxy-arp statistics command was added. |
| ARP: Displaying information about ARP direct route advertisement | The display arp route-direct advertise command was added. |
| ARP: Enabling dropping ARP requests that match FIB entries | The arp fib-miss drop command was added. |
| ARP: Testing whether an IPv4 address in a LAN is being used by sending ARP requests | The ping arp ip command was added. |
| ARP: Testing whether a MAC address exists in a specified network or to view its corresponding IPv4 address | The ping arp mac command was added. |
| DHCP: Enabling DHCP packet rate limit on an interface and setting the limit value | The dhcp rate-limit command was added. |
| DHCP: Enabling IP exhaustion event logging | The exhaustion log enable command was added. |
| DHCP: Enabling IP exhaustion notifications for an IP pool | The exhaustion trap enable command was added. |
| DHCP: Enabling SNMP notifications for the DHCP server | The snmp-agent trap enable dhcp server command was added. |
| DHCP: Specifying the IP address to be filled in sub-option 5 of Option 82 | The dhcp relay information link-selection command was added. |
| DHCP: Enabling release notification | The dhcp relay release-agent command was added. |
| DHCP: Enabling packet drop alarm logging | The following commands were added: <ul style="list-style-type: none"> • dhcp snooping alarm enable (system view) • dhcp snooping alarm enable (interface view) |
| DHCP: Setting a packet drop alarm threshold | The following commands were added: <ul style="list-style-type: none"> • dhcp snooping alarm threshold (system view) • dhcp snooping alarm threshold (interface view) |
| DHCP: Enabling the giaddr field check in DHCP requests | The dhcp snooping check giaddr command was added. |
| DHCP: Enabling client offline detection | The dhcp snooping client-detect command was |

| Feature | Command changes |
|--|--|
| | added. |
| DHCP: Enabling DHCP snooping entry exhaustion notifications | The dhcp snooping exhaustion trap enable command was added. |
| DHCP: Setting the DHCP snooping entry usage threshold | The dhcp snooping learning-num-threshold command was added. |
| DHCP: Enabling recording untrusted DHCP servers | The dhcp snooping untrusted-server-record enable command was added. |
| DHCP: Displaying statistics for dropped DHCP packets on an interface of the DHCP snooping device | The display dhcp snooping alarm packet statistics command was added. |
| DHCP: Displaying statistics about the packets exchanged between M-LAG member devices for DHCP snooping entry synchronization | The display dhcp snooping m-lag-statistics command was added. |
| DHCP: Displaying M-LAG status information | The display dhcp snooping m-lag-status statistics command was added. |
| DHCP: Clearing statistics for dropped DHCP packets on a DHCP snooping device | The reset dhcp snooping alarm packet statistics command was added. |
| DHCP: Clearing statistics about the packets exchanged between M-LAG member devices for DHCP snooping entry synchronization | The reset dhcp snooping m-lag-statistics command was added. |
| DHCP: Enabling SNMP notifications for DHCP snooping events | The snmp-agent trap enable dhcp snooping command was added. |
| IP forwarding basics: Enabling SNMP notifications for FIB events | The snmp-agent trap enable fib command was added. |
| IP forwarding basics: Displaying the FIB table usage | The display fib usage command was added. |
| IP forwarding basics: Enabling hardware forwarding for specific packets received from VXLAN tunnels | The forwarding vxlan-packet inner-protocol command was added. |
| IP performance optimization: Enabling the device to respond to broadcast echo requests | The ip icmp broadcast-echo-reply enable command was added. |
| IP performance optimization: Enabling the device to send a specific type of ICMP messages | The ip icmp send enable command was added. |
| IP performance optimization: Enabling the device to receive a specific type of ICMP messages | The ip icmp receive enable command was added. |
| IP performance optimization: Enabling SNMP notifications for TCP events | The snmp-agent trap enable tcp command was added. |
| ARP: Configuring the description for a static ARP entry | The [description text] option was added to the arp static command. |
| ARP: Configuring the description for a multiport ARP entry | The [description text] option was added to the arp |

| Feature | Command changes |
|---|---|
| | multiport command. |
| ARP: Displaying detailed information about ARP entries for a VPN instance | The verbose keyword was added to the display arp vpn-instance command. |
| ARP: Setting the preference and route tag for ARP-advertised direct routes | The preference <i>preference-value</i> and tag <i>tag-value</i> options were added to the arp route-direct advertise command. |
| DHCP: Retaining the original address in the giaddr field in relayed DHCP requests | The default-giaddr keyword was added to the dhcp relay source-address command. |
| IP forwarding basics: Specifying an IPv6 flow label for ECMP route selection and displaying the load sharing path selected for a flow | The flow-label <i>flow-label</i> option was added to the display ip load-sharing path command. |
| Tunneling: Enabling SNMP notifications for tunneling | The snmp-agent trap enable tunnel [vxlan-tunnel-status vxlan-ipv6-tunnel-status] * command was added. |
| IPv6 basics: Displaying complete IPv6 interface descriptions | The description keyword was added to the display ipv6 interface command. |
| IPv6 basics: Displaying statistics for ND proxy reply packets | The display ipv6 nd proxy statistics command was added. |
| IPv6 basics: Displaying information about ND direct route advertisement | The display ipv6 nd route-direct advertise interface <i>interface-type interface-number</i> command was added. |
| IPv6 basics: Displaying the ND table usage | The display ipv6 neighbors usage command was added. |
| IPv6 basics: Dropping the NS packets whose destination addresses already exist in the FIB table | The ipv6 nd fib-miss drop command was added. |
| IPv6 basics: Enabling logging ND entry deployment events | The ipv6 nd hardware log enable [count-limit count-limit-value] command was added. |
| IPv6 basics: Configuring a static neighbor entry | The following options were added to the ipv6 neighbor command: <ul style="list-style-type: none"> • vsi-interface <i>vsi-interface-id</i> • tunnel <i>number</i> • vsi <i>vsi-name</i> • service-instance <i>instance-id</i> |
| IPv6 basics: Setting the maximum number of probes to test the reachability of neighbors in ND entries | The ipv6 neighbor aging probe-count <i>count</i> command was added. |
| IPv6 basics: Setting the interval for testing the reachability of neighbors in ND entries. | The ipv6 neighbor aging probe-interval <i>interval</i> command was added. |
| IPv6 basics: Verifying the availability of an IPv6 address in the LAN | The ping nd ipv6 host [interface interface-type interface-number [vlan vlan-id]] [timeout timeout] [count count] command was added. |
| IPv6 basics: Obtaining the IPv6 | The ping nd mac <i>mac-address</i> { interface |

| Feature | Command changes |
|---|--|
| address of the device that uses the specified MAC address in a specific subnet | <code>interface-type interface-number ipv6 ipv6-address [vpn-instance vpn-instance-name] } [timeout timeout] [count count]</code> command was added. |
| DHCPv6: Enabling IPv6 resource exhaustion logging | The exhaustion log enable command was added. |
| DHCPv6: Enabling IPv6 resource exhaustion alarming for an IPv6 address pool | The exhaustion trap enable command was added. |
| DHCPv6: Setting the IPv6 address usage threshold for an IPv6 address pool | The ip-in-use threshold threshold-value command was added. |
| DHCPv6: Setting the prefix usage threshold for an IPv6 address pool | The pd-in-use threshold threshold-value command was added. |
| DHCPv6: Enabling SNMP notifications for the DHCPv6 server | The snmp-agent trap enable ipv6 dhcp server [address-exhaust ip-in-use pd-exhaust pd-in-use] * command was added. |
| DHCPv6: Configuring a suboption for a DHCPv6 vendor-specific option (Option 17) | The suboption suboption-code { address ipv6-address&<1-4> ascii ascii-string hex hex-string } command was added. |
| DHCPv6: Configuring a DHCPv6 vendor-specific option for an IPv6 address pool and enter the option view | The vendor-specific vendor-id command was added. |
| DHCPv6: Setting the aging timer for DUID entries on a DHCPv6 relay interface | The ipv6 dhcp relay duid aging-time seconds command was added. |
| DHCPv6: Specifying the DUID to be set in Option 37 | The ipv6 dhcp relay remote-id duid { ascii ascii-string hex hex-string } command was added. |
| DHCPv6: Enabling the DHCPv6 relay agent to add Option 37 in Relay-forward messages | The ipv6 dhcp relay remote-id enable command was added. |
| DHCPv6: Configuring the DHCPv6 relay agent to discard the DHCPv6 requests that are delivered from VXLAN tunnels | The ipv6 dhcp relay request-from-tunnel discard command was added. |
| DHCPv6: Configuring smart relay on a DHCPv6 relay interface | The ipv6 dhcp smart-relay { count count time seconds } * command was added. |
| DHCPv6: Enabling smart relay on a DHCPv6 relay interface | The ipv6 dhcp smart-relay enable command was added. |
| DHCPv6: Enabling the packet drop alarm | The ipv6 dhcp snooping alarm { relay-forward request-message } enable command was added. |
| DHCPv6: Setting a packet drop alarm threshold | The ipv6 dhcp snooping alarm { relay-forward request-message } threshold threshold command was added. |
| DHCPv6: Enabling support for the interface-ID option (also called Option 18) | The ipv6 dhcp snooping option interface-id enable command was added. |

| Feature | Command changes |
|--|---|
| DHCPv6: Enabling the lightweight DHCPv6 relay agent (LDRA) on an interface | The <code>ipv6 dhcp snooping relay-agent enable [trust]</code> command was added. |

New features: Layer 3—IP routing features

Table 4 describes the Layer 3 IP services features added in this software version.

For more information about the features, see *Layer 3—IP Routing Configuration Guide* in *HPE 5940 Switch Series Configuration Guides-Release 671x*.

For more information about the commands, see *Layer 3—IP Routing Command Reference* in *HPE 5940 Switch Series Command References-Release 671x*.

Table 4 Layer 3 IP routing features added in Release 6710

| Feature | Command changes |
|--|---|
| OSPF: Displaying OSPF neighbor relationship troubleshooting information | The <code>display ospf troubleshooting</code> command was added. |
| OSPF: Enabling OSPF isolation | The <code>isolate enable</code> command was added. |
| OSPF: Associating an OSPF interface with a track entry to adjust the cost of the interface based on the track entry state. | The <code>ospf track adjust-cost</code> command was added. |
| OSPF: Setting the maximum number of OSPF neighbor relationship troubleshooting entries | The <code>ospf troubleshooting max-number</code> command was added. |
| OSPF: Clearing OSPF neighbor relationship troubleshooting information | The <code>reset ospf troubleshooting</code> command was added. |
| OSPF: Shutting down an OSPF process | The <code>shutdown process</code> command was added. |
| IS-IS: Displaying IS-IS log information about received or sent hello packets | The <code>display isis event-log hello</code> command was added. |
| IS-IS: Displaying neighbor state change log information | The <code>display isis event-log peer</code> command was added. |
| IS-IS: Displaying information about the first three and last three packets sent by an interface | The <code>display isis interface hello-sent</code> command was added. |
| IS-IS: Displaying information about hello packets received from neighbors | The <code>display isis peer hello-received</code> command was added. |
| IS-IS: Displaying IS-IS neighbor relationship troubleshooting information | The <code>display isis troubleshooting</code> command was added. |
| IS-IS: Setting the maximum number of log entries that IS-IS | The <code>event-log size</code> command was added. |

| Feature | Command changes |
|--|--|
| can record | |
| IS-IS: Setting the maximum cost from the source node of a protected link to a PQ node | The fast-reroute remote-lfa maximum-cost command was added. |
| IS-IS: Specifying a prefix list to filter remote LFA PQ nodes | The fast-reroute remote-lfa prefix-list command was added. |
| IS-IS: Enabling IS-IS remote Loop-free Alternate (LFA) FRR | The fast-reroute remote-lfa tunnel ldp command was added. |
| IS-IS: Setting the priority for a backup path selection policy | The fast-reroute tiebreaker command was added. |
| IS-IS: Disabling remote LFA calculation on an interface | The isis fast-reroute remote-lfa disable command was added. |
| IS-IS: Setting the maximum number of IS-IS neighbor relationship troubleshooting entries | The isis troubleshooting max-number command was added. |
| IS-IS: Enabling IS-IS isolation | The isolate enable command was added. |
| IS-IS: Enabling the IS-IS multi-instance process and specify an instance ID for the process | The multi-instance enable command was added. |
| IS-IS: Clearing IS-IS GR log information | The reset isis event-log graceful-restart command was added. |
| IS-IS: Clearing IS-IS neighbor relationship troubleshooting information | The reset isis troubleshooting command was added. |
| IS-IS: Shutting down an IS-IS process | The shutdown process command was added. |
| IS-IS: Generating notifications about IS-IS address family changes and IS-IS adjacency status changes (An IS-IS address family change indicates that an address family is added or deleted.) | The adjacency-protocol-change keyword was added to the snmp-agent trap enable isis command. |
| OSPFv3: Configuring a description for an OSPFv3 process. | The description command was added. |
| OSPFv3: Enabling OSPFv3 isolation | The isolate enable command was added. |
| OSPFv3: Associating an OSPFv3 interface with a track entry to adjust the cost of the interface based on the track entry state | The ospfv3 track adjust-cost command was added. |
| OSPFv3: Shutting down an OSPFv3 process | The shutdown process command was added. |
| BGP: Support for BGP dedicated address family view | The following commands were added: <ul style="list-style-type: none"> address-family dedicated |

| Feature | Command changes |
|--|---|
| | <ul style="list-style-type: none"> <code>display bgp group dedicated</code> <code>display bgp peer dedicated</code> <code>display bgp routing-table dedicated</code> <code>display bgp update-group dedicated</code> <code>refresh bgp dedicated</code> <code>reset bgp dedicated</code> <p>The following commands are supported in BGP dedicated address family view:</p> <ul style="list-style-type: none"> <code>peer allow-as-loop</code> <code>peer enable</code> <code>peer next-hop-local</code> <code>peer reflect-client</code> <code>reflect between-clients</code> <code>reflector cluster-id</code> |
| BGP: Minimizing the priority of the routes advertised to BGP peers | <p>The following commands were added:</p> <ul style="list-style-type: none"> <code>advertise lowest-priority on-peer-up duration</code> <code>advertise lowest-priority on-startup duration</code> <code>reset bgp advertise lowest-priority</code> |
| BGP: Setting the local device IP as the next hop of each BGP ECMP route for load balancing | The ecmp-nexthop-local keyword was added to the balance command. |
| BGP: Retaining the next hop of each BGP ECMP route for load balancing | The ecmp-nexthop-unchanged keyword was added to the balance command. |
| BGP: Enabling BGP to use routes with different IGP metrics to the next hop for load balancing | The balance igp-metric-ignore command was added. |
| BGP: Configuring the time that BGP must wait for other protocols to complete GR or NSR after BGP completes GR or NSR | The bgp update-delay wait-other-protocol command was added. |
| BGP: Enabling the specified BMP server to monitor all peers in all BGP-VPN instances | The bmp server monitor all-vpn-instance command was added. |
| BGP: Enabling the specified BMP server to monitor all peers in the current instance | The bmp server monitor current-instance command was added. |
| BGP: Displaying information about BGP peers monitored by the specified BMP server for the specified BGP instance | The display bgp bmp server monitor-peer command was added. |
| BGP: Displaying BGP route dampening parameters for IBGP VPNv4 routes | The vpn4 keyword was added to the display bgp dampening parameter command. |
| BGP: Displaying BGP route information by route attribute | <ul style="list-style-type: none"> The as-path and cluster-list keywords were added to the display bgp link-state command. |

| Feature | Command changes |
|---|---|
| | <ul style="list-style-type: none"> The as-path, cluster-list, community, ext-community, <i>community-number</i><1-32>, <i>aa:nn</i><1-32>, internet, no-advertise, no-export, no-export-subconfed, rt route-target, and soo site-of-origin parameters were added to the display bgp routing-table ipv4 multicast command. The as-path and cluster-list keywords were added to the display bgp routing-table ipv4 rtfilter command. The as-path, cluster-list, community, ext-community, <i>community-number</i><1-32>, <i>aa:nn</i><1-32>, internet, no-advertise, no-export, no-export-subconfed, rt route-target, and soo site-of-origin parameters were added to the display bgp routing-table ipv4 unicast command. The as-path, cluster-list, community, ext-community, <i>community-number</i><1-32>, <i>aa:nn</i><1-32>, internet, no-advertise, no-export, no-export-subconfed, rt route-target, and soo site-of-origin parameters were added to the display bgp routing-table ipv6 multicast command. The as-path, cluster-list, community, ext-community, <i>community-number</i><1-32>, <i>aa:nn</i><1-32>, internet, no-advertise, no-export, no-export-subconfed, rt route-target, and soo site-of-origin parameters were added to the display bgp routing-table ipv6 unicast command. |
| BGP: Configuring BGP NSR | <p>The following commands were added:</p> <ul style="list-style-type: none"> display bgp non-stop-routing status non-stop-routing |
| BGP: Displaying BGP peer or peer group information | The interface , <i>link-local-address</i> , and <i>interface-type interface-number</i> parameters were added to the display bgp peer command. |
| BGP: Displaying the ORF prefix information received by a peer | The interface , <i>link-local-address</i> , and <i>interface-type interface-number</i> parameters were added to the display bgp peer received prefix-list command. |
| BGP: Manually soft-resetting BGP sessions | The interface , <i>link-local-address</i> , and <i>interface-type interface-number</i> parameters were added to the refresh bgp command. |
| BGP: Resetting BGP sessions for the specified address family | The interface , <i>link-local-address</i> , and <i>interface-type interface-number</i> parameters were added to the reset bgp command. |
| BGP: Displaying dampened BGP VPNv4 routes | The vpnvp4 keyword was added to the display bgp routing-table dampened command. |
| BGP: Displaying flap statistics about BGP VPNv4 routes | The vpnvp4 keyword was added to the display bgp routing-table flap-info command. |

| Feature | Command changes |
|--|--|
| IP routing basics: Enabling maintenance probe (MTP) | The maintenance-probe enable command was added. |
| IP routing basics/static routing/IPv6 static routing: Displaying the time of the latest next hop update | <p>The Age field was added to the command output of the following commands:</p> <ul style="list-style-type: none"> • display rib nib • display route-direct nib • display ipv6 rib nib • display ipv6 route-direct nib • display route-static nib • display ipv6 route-static nib |
| IP routing basics: Specifying an alarm threshold, in percentage, on the number of active route prefixes | The warn-threshold argument was added to the routing-table limit command. |
| IPv6 static routing: Enabling periodic sending of ND requests to the next hops of IPv6 static routes | The ipv6 route-static nd-request command was added. |
| RIP: Configuring RIPv2 keychain authentication | The keychain keyword was added to the rip authentication-mode command. |
| Policy-based routing: Specifying service chains for next hops and default next hops | <p>The service-chain path-id keyword was added to the following commands:</p> <ul style="list-style-type: none"> • apply default-next-hop • apply next-hop |
| IPv6 policy-based routing: Specifying service chains for next hops and default next hops | <p>The service-chain path-id keyword was added to the following commands:</p> <ul style="list-style-type: none"> • apply default-next-hop • apply next-hop |
| Routing policy: Removing specific AS numbers or clearing the AS_PATH attribute | The delete and clear keywords were added to the apply as-path command. |
| Routing policy: Setting a cost type for IS-IS routes or modify the MED value for BGP routes | <p>The following keywords were added to the apply cost-type command:</p> <ul style="list-style-type: none"> • inherit-link-cost: Uses the IGP link cost as the MED value of BGP routes. • internal-inc-ibgp: For IS-IS, this keyword sets the cost type a matching IS-IS route to IS-IS internal route. For BGP, this keyword sets the MED value for a matching BGP route to the IGP metric of the route's next hop. • med-plus-igp: Sets the sum of the original MED value and the IGP metric value of next hop as the MED value for a matching BGP route. |
| Routing policy: Setting a label index value | The apply label-index command was added. |
| Routing policy: Setting an SID value | The apply label-value command was added. |
| Routing policy: Matching BGP EVPN Selective Provider Multicast Service Interface (S-PMSI) routes or BGP EVPN | The bgp-evpn-s-pmsi keyword and the bgp-evpn-smet keyword were added to the if-match route-type command. |

| Feature | Command changes |
|---|---|
| Selective Multicast Ethernet Tag (SMET) routes | |
| Routing policy: Specifying an index number for an extended community list entry | The index keyword was added to the ip extcommunity-list command. |
| BGP: Displaying information about BGP peer relationship down events | The display bgp troubleshooting command was added. |
| BGP/EVPN: Filtering advertised or received BGP EVPN routes | <p>The following commands are supported in BGP EVPN address family view:</p> <ul style="list-style-type: none"> • filter-policy export • filter-policy import • peer filter-policy |
| BGP: Enabling conversational remote host route learning | The forwarding-conversational-learning command was added. |
| BGP: Configuring peers that are created through interfaces | <ul style="list-style-type: none"> • interface-peer additional-paths • interface-peer advertise additional-paths best • interface-peer advertise origin-as-validation • interface-peer advertise-community • interface-peer advertise-ext-community • interface-peer advertise-policy exist-policy • interface-peer advertise-policy non-exist-policy • interface-peer allow-as-loop • interface-peer as-number • interface-peer as-path-acl • interface-peer bfd • interface-peer bmp server • interface-peer capability-advertise orf prefix-list • interface-peer default-route-advertise • interface-peer description • interface-peer dscp • interface-peer enable • interface-peer fake-as • interface-peer filter-policy • interface-peer graceful-restart timer restart extra • interface-peer group • interface-peer ignore • interface-peer ignore-first-as • interface-peer ignore-originatorid • interface-peer keep-all-routes • interface-peer keychain • interface-peer log-change |

| Feature | Command changes |
|--|---|
| | <ul style="list-style-type: none"> <code>interface-peer low-memory-exempt</code> <code>interface-peer next-hop-local</code> <code>interface-peer nexthop-recursive-policy disable</code> <code>interface-peer password</code> <code>interface-peer preferred-value</code> <code>interface-peer prefix-list</code> <code>interface-peer public-as-only</code> <code>interface-peer reflect-client</code> <code>interface-peer route-limit</code> <code>interface-peer route-policy</code> <code>interface-peer route-update-interval</code> <code>interface-peer soo</code> <code>interface-peer substitute-as</code> <code>interface-peer timer</code> <code>interface-peer timer connect-retry</code> |
| BGP: Configuring BGP to advertise the BGP RPKI validation state to a peer or peer group | The peer advertise origin-as-validation command is supported in BGP VPNv4 address family view and BGP VPNv6 address family view. |
| BGP: Enabling BFD for the link to a BGP peer or peer group | The echo keyword was added to the peer bfd command. |
| BGP: Removing private AS numbers in BGP updates sent to an EBGP peer or peer group | The force , limited , replace , and include-peer-as keywords were added to the peer public-as-only command. |
| BGP: Configuring BGP maintenance features | The shutdown process command and the isolate enable command were added. |
| BGP: Specifying the longest match principle for BGP next hop recursion | The nexthop recursive-lookup longest-match command was added. |
| BGP: Configuring the BMP client to send peer down notifications with mode flags to the BMP server | The pd-monitor-mode enable command was added. |
| BGP: Configuring the BMP client to send peer up notifications with mode flags to the BMP server | The pu-monitor-mode enable command was added. |
| BGP: Enabling BGP to check the first AS number of EBGP routes (BGP will not advertise an EBGP route to EBGP peers whose AS number is the route's first AS number.) | The peer-as-check enable command was added. |
| BGP: Enabling BGP to send routes exchanged with the specified monitored peer or peer group to the BMP server. | The peer route-mode command was added. |
| BGP: Associating a BGP peer or peer group with a track entry (BGP can adjust the priority of | The peer route-priority-track command was added. |

| Feature | Command changes |
|--|---|
| routes received from the peer or peer group based on the track entry state.) | |
| BGP: Configuring the device as a route server and specifying a peer or peer group as a client | The peer route-server-client command was added. |
| BGP: Setting the TCP maximum segment size (MSS) for a peer or peer group | The peer tcp-mss command was added. |
| BGP: Enabling peer unreachability detection | The peer tracking command was added. |
| BGP: Using control-mode BFD to detect the connectivity to the next hop of the primary route | The ctrl keyword was added to the primary-path-detect bfd command. |
| BGP: Resetting the connection to a BMP server and clear statistics information | The reset bgp bmp server command was added. |
| BGP: Clearing BGP route dampening information and release suppressed BGP VPNv4 routes | The vpn4 keyword was added to the reset bgp dampening command. |
| BGP: Clearing flap statistics for BGP VPNv4 routes | The vpn4 keyword was added to the reset bgp flap-info command. |
| BGP: Enabling BGP to send routes received from all the monitored peers and peer groups to the BMP server | The route-mode adj-rib-in command was added. |
| BGP: Configuring BGP to send routes advertised to the monitored peer or peer group to the BMP server | The pre-policy , post-policy , and both keywords were added to the route-mode adj-rib-out command. |
| BGP: Enabling BGP route advertisement delay | The route-update-delay command was added. |
| BGP: Configuring an IP address and port number for a BMP server | The ipv6-address argument was added to the server command. |
| BGP: Specifying the authentication mode and key for BGP to establish TCP connections to the BMP server | The server password command was added. |
| BGP: Specifying the source address of TCP connections to the BMP server | The server source-address command was added. |
| BGP: New BGP notification types | <p>The following keywords were added to the snmp-agent trap enable bgp command:</p> <ul style="list-style-type: none"> • peer-addrfamily-routeexceed • peer-addrfamily-routeexceed-clear • peer-addrfamily-routethreshold-clear • peer-addrfamily-routethreshold-exceed • peer-backward-transition |

| Feature | Command changes |
|--|--|
| | <ul style="list-style-type: none"> • peer-established • peer-routeexceed • peer-routeexceed-clear • peer-routethreshold-clear • peer-routethreshold-exceed |
| BGP: Enabling fast host route update upon user migration | The user-move fast-update command was added. |
| DCN features | All DCN commands were newly added. |

New features: IP multicast features

[Table 5](#) describes the IP multicast features added in this software version.

For more information about the features, see *IP Multicast Configuration Guide* in *HPE 5940 Switch Series Configuration Guides-Release 671x*.

For more information about the commands, see *IP Multicast Command Reference* in *HPE 5940 Switch Series Command References-Release 671x*.

Table 5 IP multicast features added in Release 6710

| Feature | Command changes |
|---|---|
| IGMP snooping/MLD snooping: Preventing the device from forwarding multicast data to router ports in a VLAN | <p>The following commands were added:</p> <ul style="list-style-type: none"> • igmp-snooping router-port-discard • mld-snooping router-port-discard |
| IGMP snooping/MLD snooping: Enabling the device to send IGMP/MLD general queries upon a port state change | <p>The following commands were added:</p> <ul style="list-style-type: none"> • igmp-snooping triggered-query enable • mld-snooping triggered-query enable |
| IGMP snooping/MLD snooping: Configuring a multicast access control policy | <p>The following commands were added:</p> <ul style="list-style-type: none"> • igmp-snooping access-policy • mld-snooping access-policy |
| Multicast routing and forwarding: Setting the maximum number of copied multicast packets during software forwarding | The multicast cpu-forwarding max-copy-count command was added. |
| Multicast routing and forwarding/IPv6 multicast routing and forwarding: Configuring multicast load splitting | <p>The following parameters were added to the load-splitting (IPv6 MRIB view) and load-splitting (MRIB view) commands:</p> <ul style="list-style-type: none"> • balance-ecmp • balance-ucmp • ecmp • ucmp |
| IGMP: Specifying an ACL by its name when configuring a filtering rule in IGMP | The name ipv4-acl-name option was added to the igmp group-policy and ssm-mapping commands. |

| Feature | Command changes |
|---|--|
| IGMP: Displaying IGMP multicast groups, IGMP proxy-maintained multicast groups, and IGMP proxy-maintained multicast routing entries learned by EVPN | <p>The evpn keyword was added to the following commands:</p> <ul style="list-style-type: none"> display igmp group display igmp proxy group display igmp proxy routing-table |
| PIM: Specifying an ACL by its name when configuring a filtering rule in PIM | <p>The name <i>ipv4-acl-name</i> option was added to the following commands:</p> <ul style="list-style-type: none"> c-rp pim neighbor-policy register-policy source-policy ssm-policy static-rp |
| PIM/IPv6 PIM: Setting the timer for C-BSRs to wait BSMs from the BSR | <p>The following commands were added:</p> <ul style="list-style-type: none"> c-bsr holdtime (PIM view) c-bsr holdtime (IPv6 PIM view) |
| PIM/IPv6 PIM: Setting the interval for C-BSRs to send BSMs | <p>The following commands were added:</p> <ul style="list-style-type: none"> c-bsr interval (PIM view) c-bsr interval (IPv6 PIM view) |
| PIM/IPv6 PIM: Setting the delay timer for DR election | <p>The following commands were added:</p> <ul style="list-style-type: none"> pim timer dr-elect-delay ipv6 pim timer dr-elect-delay |
| PIM/IPv6 PIM: Enabling the device to send hello messages with a different Generation ID upon a port state change | <p>The following commands were added:</p> <ul style="list-style-type: none"> pim triggered-hello enable ipv6 pim triggered-hello enable |
| PIM/IPv6 PIM: Setting the interval for C-BSRs to send BSMs | <p>The following commands were added:</p> <ul style="list-style-type: none"> c-bsr interval (PIM view) c-bsr interval (IPv6 PIM view) |
| PIM/IPv6 PIM: Configuring a PIM/IPv6 PIM join policy to filter joined multicast sources and groups in PIM join or prune messages on an interface | <p>The following commands were added:</p> <ul style="list-style-type: none"> pim join-policy ipv6 pim join-policy |
| PIM/IPv6 PIM: Configuring an interface as a DR interface | <p>The following commands were added:</p> <ul style="list-style-type: none"> pim distributed-dr ipv6 pim distributed-dr |
| PIM/IPv6 PIM: Enabling SNMP notifications | <p>The interface-election and rp-mapping-change parameters were added to the following commands:</p> <ul style="list-style-type: none"> snmp-agent trap enable pim snmp-agent trap enable pim6 |
| Multicast VPN: Displaying BGP MDT routing information | <p>The as-path and cluster-list parameters were added to the display bgp routing-table ipv4 mdt command.</p> |
| Multicast VPN: Displaying BGP IPv4 MVPN routing information | <p>The following parameters were added to the display bgp routing-table ipv4 mvpn</p> |

| Feature | Command changes |
|--|--|
| | command: <ul style="list-style-type: none"> • as-path • cluster-list • ext-community • rt <i>route-target</i> • whole-match |
| IPv6 multicast routing and forwarding: Deleting all static IPv6 multicast routes | The delete ipv6 rpf-route-static command was added. |
| IPv6 multicast routing and forwarding: Displaying static IPv6 multicast routing entries | The display ipv6 multicast routing-table static command was added. |
| IPv6 multicast routing and forwarding: Configuring a static IPv6 multicast route | The ipv6 rpf-route-static command was added. |
| IPv6 multicast routing and forwarding: Setting the maximum number of copied multicast packets during software forwarding | The ipv6 multicast cpu-forwarding max-copy-count command was added. |
| MLD: Specifying an IPv6 ACL by its name when configuring an IPv6 multicast group policy or MLD SSM mapping | The name <i>ipv6-acl-name</i> option was added to the mld group-policy and ssm-mapping commands. |
| IPv6 PIM: Specifying an IPv6 ACL by its name when configuring a filtering rule in IPv6 PIM | The name <i>ipv6-acl-name</i> option was added to the following commands: <ul style="list-style-type: none"> • static-rp • c-rp • register-policy • ssm-policy • source-policy • ipv6 pim neighbor-policy |
| IPv6 PIM: Displaying information about the IPv6 PIM routing entries for MVPN extranet | The extranet parameter was added to the display ipv6 pim routing-table command. |

New features: MPLS features

Table 6 describes the MPLS features added in this software version.

For more information about the features, see *MPLS Configuration Guide* in *HPE 5940 Switch Series Configuration Guides-Release 671x*.

For more information about the commands, see *MPLS Command Reference* in *HPE 5940 Switch Series Command References-Release 671x*.

Table 6 MPLS features added in Release 6710

| Feature | Command changes |
|---|--|
| MPLS L2VPN: Displaying L2VPN PW state machine information | The display l2vpn pw state-machine command was added. |
| MPLS L2VPN: Specifying the PW | The independent keyword was added to the |

| Feature | Command changes |
|--|--|
| redundancy mode as independent | pw-redundancy command. |
| VPLS: Specifying the PW redundancy mode as independent | The independent keyword was added to the pw-redundancy command. |
| SR-MPLS: Assigning an IS-IS adjacency SID to an IS-IS adjacency | The isis adjacency-sid command was added. |
| SR-MPLS: Displaying IS-IS SR prefix-SID mappings | The display isis segment-routing prefix-sid-map command was added. |
| SR-MPLS: Configuring an IS-IS prefix SID for an IS-IS multi-instance process | The process-id process-id option was added to the isis prefix-sid command. |
| SR-MPLS: Setting the P-flag bit of the prefix SID to 1 so the penultimate hop does not pop out the SID | The no-php keyword was added to the isis prefix-sid command. |
| SR-MPLS: Enabling SR-MPLS for the BGP IPv4 unicast address family | The segment-routing mpls command was added to BGP IPv4 unicast address family view. |
| SR-MPLS: Displaying BGP SR label ranges | The display bgp segment-routing label-range command was added. |
| SR-MPLS: Enabling BGP EPE for a peer or peer group and specifying a routing policy to assign a label (SID) to the peer or peer group | The peer egress-engineering command was added. |
| SR-MPLS: Displaying BGP EPE adjacency segment information | The display bgp egress-engineering ipv4 command was added. |
| SR-MPLS: Displaying prefix-SID mappings | The display segment-routing mapping-server prefix-sid-map command was added. |
| SR-MPLS: Configuring a prefix-SID mapping | The mapping-server prefix-sid-map command was added. |
| SR-MPLS: Enabling segment routing and enter segment routing view | The segment-routing command was added. |
| SR-MPLS: Enabling advertisement of locally configured prefix-SID mappings | The segment-routing mapping-server advertise-local command was added. |
| SR-MPLS: Enabling reception of prefix-SID mappings | The segment-routing mapping-server receive command was added. |
| SR-MPLS: Enabling FRR microloop avoidance | The fast-reroute microloop-avoidance enable command was added. |
| SR-MPLS: Setting the FRR microloop avoidance RIB-update-delay time | The fast-reroute microloop-avoidance rib-update-delay command was added. |
| SR-MPLS: Enabling Topology-Independent Loop-Free Alternate Fast Re-Route (TI-LFA FRR) | The fast-reroute ti-lfa command was added. |
| SR-MPLS: Disabling an IS-IS interface from participating in TI-LFA calculation | The isis fast-reroute ti-lfa disable command was added. |
| SR-MPLS: Enabling SR microloop avoidance | The segment-routing microloop-avoidance enable command was added. |
| SR-MPLS: Setting the SR microloop avoidance RIB-update-delay time | The segment-routing microloop-avoidance rib-update-delay command was added. |

| Feature | Command changes |
|---|---|
| LDP: Enabling the device to automatically establish targeted LDP sessions on receiving LDP Targeted Hellos | The accept target-hello { all prefix-list prefix-list-name } command was added. |
| LDP: Displaying error TCP packets received by the LDP session | The display mpls ldp error-packet command was added. |
| LDP: Display information about LDP session interruptions | The display mpls ldp troubleshooting command was added. |
| LDP: Configuring a label mapping propagation policy | propagate mapping { none prefix-list prefix-list-name } |
| MPLS TE: Adding or modifying a label node in an explicit path, and configure the attributes of the label node | The nextsid [index index-number] label label-value type { adjacency prefix } command was added. |
| MPLS OAM: Verifying the MPLS LSP of the specified outgoing labels | The ping mpls [-a source-ip -c count -exp exp-value -h ttl-value -m wait-time -r reply-mode -rtos tos-value -s packet-size -t time-out -v] * out-labels out-label-value&<1-3> interface interface-type interface-number [nexthop nexthop-address] command was added. |
| MPLS OAM: Tracing MPLS LSPs of the specified outgoing labels from the ingress node to the egress node of the LSPs | The tracert mpls [-a source-ip -exp exp-value -h ttl-value -r reply-mode -rtos tos-value -t time-out -v fec-check] * out-labels out-label-value&<1-3> interface interface-type interface-number [nexthop nexthop-address] command was added. |
| MPLS L3VPN: Configuring BGP VPNv4 route dampening | The dampening ibgp command was added. |

New features: ACL and QoS features

Table 7 describes the ACL and QoS features added in this software version.

For more information about the features, see *ACL and QoS Configuration Guide* in *HPE 5940 Switch Series Configuration Guides-Release 671x*.

For more information about the commands, see *ACL and QoS Command Reference* in *HPE 5940 Switch Series Command References-Release 671x*.

Table 7 ACL and QoS features in Release 6710

| Feature | Command changes |
|---|--|
| ACL: Ignoring the permit flag added by packet filtering | The packet-filter permit-flag ignore command was added. |
| ACL: Matching object groups in an ACL rule | <p>The object-group port-group-name option was added to the following commands:</p> <ul style="list-style-type: none"> rule (IPv4 advanced ACL view) rule (IPv4 basic ACL view) |

| Feature | Command changes |
|--|---|
| | <ul style="list-style-type: none"> rule (IPv6 advanced ACL view) rule (IPv6 basic ACL view) |
| ACL: Matching microsegments in an advanced ACL rule | <p>The microsegment <i>microsegment-id</i> [mask-length <i>mask-length</i>] option was added to the following commands:</p> <ul style="list-style-type: none"> rule (IPv4 advanced ACL view) rule (IPv6 advanced ACL view) |
| ACL: Matching inner TCP flags of VXLAN packets in an advanced ACL rule | <p>The following option was added to the rule (IPv4 advanced ACL view) and rule (IPv6 advanced ACL view) commands:</p> <pre>{ inner-ack <i>inner-ack-value</i> inner-fin <i>inner-fin-value</i> inner-psh <i>inner-psh-value</i> inner-rst <i>inner-rst-value</i> inner-syn <i>inner-syn-value</i> inner-urg <i>inner-urg-value</i> } *</pre> |
| ACL: Matching the inner IP precedence in an IPv4 advanced ACL rule | <p>The inner-precedence <i>inner-precedence</i> option was added to the rule (IPv4 advanced ACL view) command.</p> |
| ACL: Matching the inner ToS in an IPv4 advanced ACL rule | <p>The inner-tos <i>inner-tos</i> option was added to the rule (IPv4 advanced ACL view) command.</p> |
| ACL: Matching the inner ECN flag in an IPv4 advanced ACL rule | <p>The inner-ecn <i>inner-ecn</i> option was added to the rule (IPv4 advanced ACL view) command.</p> |
| ACL: Matching the inner DSCP value in an IPv4 advanced ACL rule | <p>The inner-dscp <i>inner-dscp</i> option was added to the rule (IPv4 advanced ACL view) command.</p> |
| ACL: Matching the inner header information of VXLAN packets in an IPv6 advanced ACL rule | <p>The following command was added:</p> <pre>rule [<i>rule-id</i>] { deny permit } vxlan [vxlan-id <i>vxlan-id</i>] inner-protocol <i>inner-protocol</i> [counting inner-destination { <i>dest-address</i> <i>dest-prefix</i> <i>dest-address/dest-prefix</i> any } inner-destination-port <i>operator</i> <i>port1</i> [<i>port2</i>] { { inner-ack <i>inner-ack-value</i> inner-fin <i>inner-fin-value</i> inner-psh <i>inner-psh-value</i> inner-rst <i>inner-rst-value</i> inner-syn <i>inner-syn-value</i> inner-urg <i>inner-urg-value</i> } * inner-established } inner-source { <i>source-address</i> <i>source-prefix</i> <i>source-address/source-prefix</i> any } inner-source-port <i>operator</i> <i>port1</i> [<i>port2</i>] { inner-dscp <i>inner-dscp</i> inner-ecn <i>inner-ecn</i> } * logging time-range <i>time-range-name</i>] *</pre> |
| QoS: Mirroring-type QoS policy | <p>The mirroring keyword was added to the following commands:</p> <ul style="list-style-type: none"> display qos policy interface |

| Feature | Command changes |
|---|--|
| | <ul style="list-style-type: none"> <code>display qos policy global</code> <code>display qos policy user-defined</code> <code>display qos policy l2vpn-ac</code> <code>display qos policy diagnosis global</code> <code>display qos policy diagnosis interface</code> <code>qos apply policy</code> (Ethernet service instance view, interface view, control plane view, VSI view) <code>qos apply policy global</code> <code>qos policy</code> <code>reset qos policy global</code> |
| QoS: Clearing the QoS policies applied to Ethernet service instances | The <code>reset qos policy l2vpn-ac</code> command was added. |
| QoS: Configuring the global priority trust mode for VXLAN packets | The <code>tunnel-dscp</code> keyword was added to the <code>qos trust</code> command. |
| QoS: Displaying diagnostic information about the QoS policies applied to Ethernet service instances | The <code>outbound</code> keyword was added to the <code>display qos policy diagnosis l2vpn-ac</code> command. |
| QoS: Displaying the QoS policies applied to Ethernet service instances | The <code>outbound</code> keyword was added to the <code>display qos policy l2vpn-ac</code> command. |

New features: Security features

[Table 8](#) describes the security features added in this software version.

For more information about the features, see *Security Configuration Guide* in *HPE 5940 Switch Series Configuration Guides-Release 671x*.

For more information about the commands, see *Security Command Reference* in *HPE 5940 Switch Series Command References-Release 671x*.

Table 8 Security features added in Release 6710

| Feature | Command changes |
|---|---|
| AAA: Configuring the authentication failure policy for users in an ISP domain | The <code>authen-fail</code> command was added. |
| AAA: Specifying an action to take on users in the critical domain when a RADIUS server becomes available | The <code>authen-radius-recover</code> command was added. |
| AAA: Specifying a critical domain to accommodate users that access the ISP domain when all RADIUS servers are unavailable | The <code>authen-radius-unavailable online domain</code> command was added. |
| AAA: Configuring authorization attributes for users in an ISP domain | The <code>vlan</code> and <code>vsi</code> keywords were added to the <code>authorization-attribute</code> command. |
| AAA: Enabling temporary redirect | The <code>redirect move-temporarily enable</code> command was added. |
| AAA: Configuring authorization attributes for a network access user or user group | The <code>url</code> keyword was added to the |

| Feature | Command changes |
|--|--|
| | authorization-attribute command. |
| AAA: Configuring local guests | <p>The guest keyword was added to the local-user and display local-user commands.</p> <p>The following commands were added:</p> <ul style="list-style-type: none"> • company • display local-guest waiting-approval • email • full-name • local-guest email format • local-guest email sender • local-guest email smtp-server • local-guest generate • local-guest manager-email • local-guest send-email • local-guest timer • phone • reset local-guest waiting-approval • sponsor-department • sponsor-email • sponsor-full-name • validity-datetime |
| AAA: Exporting local guest account information to a .csv file in the specified path | The local-user-export command was added. |
| AAA: Importing local guest account information from a .csv file in the specified path to the device to create local guests based on the imported information | The local-user-import command was added. |
| AAA: Configuring the format of the RADIUS NAS-Port attribute | The attribute 5 format command was added. |
| AAA: Configuring the MAC address format for the RADIUS Called-Station-Id attribute | The attribute 30 mac-format command was added. |
| AAA: Configuring the MAC address format for the RADIUS Calling-Station-Id attribute | The one keyword was added to the attribute 31 mac-format command. |
| AAA: Configuring the format of RADIUS attribute 87 | The attribute 87 format command was added. |
| AAA: Including subattribute 218 of vendor 25506 in outgoing RADIUS packets | The include-attribute 218 vendor-id 25506 command was added. |
| AAA: Specifying a source IP address for outgoing RADIUS packets | The nas-ip command was added. |
| AAA: Configuring the device to preferentially process RADIUS authentication requests | The radius authentication-request first command was added. |
| AAA: Setting the interval at which the device detects the status of RADIUS authentication | The radius-server authen-state-check interval command was added. |

| Feature | Command changes |
|--|--|
| servers | |
| AAA: Displaying HWTACACS packet statistics | The display hwtacacs statistics command was added. |
| AAA: Setting the DSCP priority for RADIUS packets | The hwtacacs dscp command was added. |
| 802.1X: Displaying online 802.1X user information | The m-lag [local peer] and online-type { auth-fail-domain critical-domain preauth-domain success } parameters were added to the display dot1x connection command. |
| 802.1X: Displaying unknown source MAC addresses in the unicast-trigger quiet period | The display dot1x unicast-trigger quiet-mac command was added. |
| 802.1X: Enabling generation of dynamic IPv4SG or IPv6SG binding entries for 802.1X authenticated users | The dot1x { ip-verify-source ipv6-verify-source } enable command was added. |
| 802.1X: Configuring 802.1X | <p>The following commands were added to Layer 2 aggregate interface view:</p> <ul style="list-style-type: none"> • dot1x after-mac-auth max-attempt • dot1x auth-fail eapol • dot1x auth-fail vlan • dot1x critical eapol • dot1x critical vlan • dot1x critical-voice-vlan • dot1x eap-tls-fragment to-server • dot1x eapol untag • dot1x guest-vlan • dot1x guest-vlan-delay • dot1x guest-vsi-delay • dot1x handshake • dot1x handshake reply enable • dot1x handshake secure • dot1x mac-binding • dot1x mac-binding enable • dot1x mandatory-domain • dot1x multicast-trigger • dot1x offline-detect enable • dot1x port-control • dot1x port-method • dot1x re-authenticate • dot1x re-authenticate manual • dot1x re-authenticate server-unreachable keep-online • dot1x server-recovery online-user-sync • dot1x timer reauth-period |

| Feature | Command changes |
|---|--|
| | <ul style="list-style-type: none"> dot1x unauthenticated-user aging enable dot1x unicast-trigger dot1x user-ip freeze |
| 802.1X: Enabling online 802.1X users to escape from offline detection and stay online when no reachable RADIUS authentication servers are available | The dot1x auth-server-unavailable escape command was added. |
| 802.1X: Discarding duplicate 802.1X EAPOL-Start requests | The dot1x duplicate-eapol-start discard command was added. |
| 802.1X: Configuring the redirect URL if users will use Web browsers to access the network | The secondary and track track-entry-number parameters were added to the dot1x ead-assistant url command. |
| 802.1X: Setting the maximum number of concurrent 802.1X users on a port | The preauth-domain and auth-fail-domain keywords were added to the dot1x max-user command. |
| 802.1X: Setting an 802.1X timer | The unicast-trigger quiet-period quiet-period-value option was added to the dot1x timer command. |
| 802.1X: Logging off 802.1X users | The online-type { auth-fail-domain critical-domain preauth-domain success } parameters were added to the reset dot1x access-user command. |
| 802.1X: Removing the records of unknown source MAC addresses in the unicast-trigger quiet period | The reset dot1x unicast-trigger quiet-mac command was added. |
| MAC authentication: Displaying information about online MAC authentication users | The m-lag [local peer] and online-type { auth-fail-domain critical-domain preauth-domain success url-unavailable-domain } parameters were added to the display mac-authentication connection command. |
| MAC authentication: Displaying MAC authentication user recovery profiles | The display mac-authentication user-recovery-profile command was added. |
| MAC authentication: Configuring the username and password for accessing the RESTful server | The login-name command was added. |
| MAC authentication: Enabling online MAC authentication users to escape from offline detection and stay online when no reachable RADIUS authentication servers are available | The mac-authentication auth-server-unavailable escape command was added. |
| MAC authentication: Enabling automatic MAC authentication user recovery | The mac-authentication auto-recover-user command was added. |
| MAC authentication: Configuring MAC authentication | <p>The following commands were added to Layer 2 aggregate interface view:</p> <ul style="list-style-type: none"> mac-authentication critical vlan mac-authentication |

| Feature | Command changes |
|---|--|
| | critical-voice-vlan <ul style="list-style-type: none"> • mac-authentication domain • mac-authentication guest-vlan • mac-authentication guest-vlan auth-period • mac-authentication offline-detect enable • mac-authentication parallel-with-dot1x • mac-authentication re-authenticate • mac-authentication re-authenticate server-unreachable keep-online • mac-authentication server-recovery online-user-sync • mac-authentication timer • mac-authentication unauthenticated-user aging enable |
| MAC authentication: Specifying the MAC authentication critical VLAN on the port | The url-user-logoff keyword was added to the mac-authentication critical vlan command. |
| MAC authentication: Enabling guest VLAN reauthentication in MAC authentication | The mac-authentication guest-vlan re-authenticate command was added. |
| MAC authentication: Enabling guest VSI reauthentication in MAC authentication | The mac-authentication guest-vsi re-authenticate command was added. |
| MAC authentication: Configuring a username and password for MAC authentication users in a MAC address range | The mac-authentication mac-range-account command was added. |
| MAC authentication: Setting the maximum number of concurrent MAC authentication users on a port | The preauth-domain and auth-fail-domain keywords were added to the mac-authentication max-user command. |
| MAC authentication: Manually recovering MAC authentication users | The mac-authentication recover-user command was added. |
| MAC authentication: Associating a redirect URL for MAC authentication users with a track entry | The mac-authentication redirect-url command was added. |
| MAC authentication: Configuring a MAC authentication timer | The temporary-user-aging aging-time-value option was added to the mac-authentication timer command. |
| MAC authentication: Configuring the global user account policy for all MAC authentication users | The separator colon keywords were added to the mac-authentication user-name-format command. |
| MAC authentication: Creating a profile for MAC authentication user recovery and enter its view | The mac-authentication user-recovery-profile command was added. |

| Feature | Command changes |
|--|--|
| MAC authentication: Configuring Web proxy ports for URL redirection in MAC authentication | The mac-authentication web-proxy command was added. |
| MAC authentication: Configuring the NAS IP address used by the device to communicate with the RESTful server | The nas-ip command was added. |
| MAC authentication: Logging off MAC authentication users | The online-type { auth-fail-domain critical-domain preauth-domain success url-unavailable-domain } parameters were added to the reset mac-authentication access-user command. |
| MAC authentication: Configuring the IP address and port number of the RESTful server | The server-address command was added. |
| Portal authentication: Setting the aging time for MAC-trigger entries | The aging-time command was added. |
| Portal authentication: Setting the timeout the device waits for portal authentication to complete after receiving the MAC binding query response | The authentication-timeout command was added. |
| Portal authentication: Setting the maximum number of attempts and the interval for sending MAC binding queries to the MAC binding server | The binding-retry command was added. |
| Portal authentication: Enabling cloud MAC-trigger authentication | The cloud-binding enable command was added. |
| Portal authentication: Specifying the URL of the cloud portal authentication server | The cloud-server url command was added. |
| Portal authentication: Displaying information about MAC binding servers | The display portal mac-trigger-server command was added. |
| Portal authentication: Displaying packet statistics for portal authentication servers | The extend-auth-server cloud keywords were added to the display portal packet statistics command. |
| Portal authentication: Displaying session information for portal users or portal-based Web authentication users | The display portal session user-type command was added. |
| Portal authentication: Displaying portal user information | The following parameters were added to the display portal user command: <ul style="list-style-type: none"> auth-type { cloud local normal } mac mac-address username username brief |
| Portal authentication: Setting the free-traffic threshold | The free-traffic threshold command was added. |
| Portal authentication: Specifying the IPv4 address of a MAC binding server | The ip (MAC binding server view) command was added. |
| Portal authentication: Specifying the IPv6 address of a MAC binding server | The ipv6 (MAC binding server view) command was added. |
| Portal authentication: Binding an SSID or endpoint | The logon-page bind command was added. |

| Feature | Command changes |
|--|--|
| name to an authentication page file | |
| Portal authentication: Configuring the NAS-Port-Type attribute carried in outgoing RADIUS requests on the interface | The nas-port-type command was added. |
| Portal authentication: Setting the maximum number of portal users allowed on an interface | The max-user keyword was added to the portal { ipv4-max-user ipv6-max-user max-user } command. |
| Portal authentication: Obtaining user access information from ARP or ND entries | The portal access-info trust command was added. |
| Portal authentication: Specifying a MAC binding server on an interface | The portal [ipv6] apply mac-trigger-server command was added. |
| Portal authentication: Specifying a portal Web server for redirect of HTTP requests sent by unauthenticated portal users | The secondary keyword was added to the portal apply web-server command. |
| Portal authentication: Logging out online portal users | The auth-type { cloud local normal } and mac mac-address parameters were added to the portal delete-user command. |
| Portal authentication: Setting the authentication load sharing mode for portal users on M-LAG interfaces | The portal m-lag load-sharing-mode command was added. |
| Portal authentication: Setting the traffic backup interval and threshold for portal users on M-LAG interfaces | The portal m-lag traffic backup command was added. |
| Portal authentication: Enabling the portal fail-permit feature for portal Web servers | The portal fail-permit web-server command was added. |
| Portal authentication: Setting the user synchronization interval for portal authentication using OAuth | The portal oauth user-sync interval command was added. |
| Portal authentication: Clearing packet statistics for portal authentication servers. | The extend-auth-server cloud keywords were added to the reset portal packet statistics command. |
| Portal authentication: Specifying the type of a portal authentication server or portal Web server | The oauth keyword was added to the server-type command. |
| Portal authentication: Specifying the type of a MAC binding server | The server-type command was added. |
| Portal authentication: Configuring the parameters to be carried in the URL when the device redirects it to users | The format section { 1 3 6 } { lowercase uppercase } keywords were added to the url-parameter command. |
| Portal authentication: Specifying the version of the portal protocol | The version command was added. |
| Web authentication: Displaying Web authentication user information | The m-lag [local peer] keywords were added to the display web-auth user command. |
| Web authentication: Configuring the Web authentication server to support the Oasis | The server-type oauth command was added. |

| Feature | Command changes |
|--|---|
| platform | |
| Web authentication: Logging off Web authentication users | The reset web-auth access-user command was added. |
| Web authentication: Configuring the redirection URL of the Web server for Web authentication | The track <i>track-entry-number</i> option was added to the url command. |
| Web authentication: Configuring the unescaped special characters in the Web authentication URL redirected to users | The url-unescape-chars command was added. |
| Web authentication: Applying a portal MAC binding server for Web authentication | The web-auth apply portal mac-trigger-server command was added. |
| Web authentication: Creating a remote Web server for Web authentication | <p>The following commands were added:</p> <ul style="list-style-type: none"> • web-auth remote server <i>server-name</i> • ip (Web authentication remote Web server view) • ipv6 (Web authentication remote Web server view) • server-type oauth • url • url-parameter • url-unescape-chars |
| Web authentication: Configuring Web authentication | <p>The following commands were added to Layer 2 aggregate interface view:</p> <ul style="list-style-type: none"> • web-auth auth-fail vlan • web-auth domain • web-auth enable • web-auth offline-detect |
| Web authentication: Enabling online detection bypass for Web authentication users | The web-auth auth-server-unavailable escape command was added. |
| Web authentication: Enabling Web authentication | The secondary-server <i>secondary-server-name</i> option was added to the web-auth enable command. |
| Web authentication: Configuring a Web authentication-free host name | The web-auth free-host command was added. |
| Web authentication: Configuring Web authentication multi-VLAN mode | The web-auth host-mode multi-vlan command was added. |
| Web authentication: Setting the maximum number of Web authentication users | The preauth-domain and auth-fail-domain keywords were added to the web-auth max-user command. |
| Web authentication: Adding the port number of a Web proxy server | The https keyword was added to the web-auth proxy port command. |
| Web authentication: Configuring the aging timer for temporary MAC address entries | The web-auth timer temp-entry-aging command was added. |
| Port security: Displaying entries for online port | The display port-security |

| Feature | Command changes |
|---|--|
| security access users | access-user command was added. |
| Port security: Displaying port security statistics | The display port-security statistics command was added. |
| Port security: Configuring port security | <p>The following commands were added to Layer 2 aggregate interface view:</p> <ul style="list-style-type: none"> • port-security authentication open • port-security authorization ignore • port-security escape critical-vsi • port-security intrusion-mode • port-security mac-address aging-type inactivity • port-security mac-address dynamic • port-security mac-address security • port-security mac-limit • port-security max-mac-count • port-security nas-id-profile • port-security port-mode |
| Port security: Configuring an authentication load sharing mode for users attached to M-LAG interfaces | The port-security m-lag load-sharing-mode command was added. |
| Port security: Enabling VLAN check bypass for users moving to the port from other ports | The port-security mac-move bypass-vlan-check command was added. |
| Port security: Enabling MAC move | The port and vlan keywords were added to the port-security mac-move permit command. |
| Port security: Setting the port security mode of a port | The mac-and-userlogin-secure-ext keyword was added to the port-security port-mode command. |
| Port security: Specifying a preauthentication domain for port security users on a port | The port-security pre-auth domain command was added. |
| Port security: Configuring static users for port access authentication | <p>The following commands were added:</p> <ul style="list-style-type: none"> • display port-security static-user • display port-security static-user connection • port-security static-user • port-security static-user match-mac acl • port-security static-user max-user |

| Feature | Command changes |
|--|---|
| | <ul style="list-style-type: none"> <code>port-security static-user password</code> <code>port-security static-user timer detect-period</code> <code>port-security static-user timer offline-detect</code> <code>port-security static-user update-ip enable</code> <code>port-security static-user user-name-format</code> <code>port-security static-user user-name-format mac-address</code> <code>reset port-security static-user</code> |
| Port security: Setting port security timers | The <code>port-security timer</code> command was added. |
| Port security: Configuring the trigger order for authentication methods on the port as MAC authentication, 802.1X authentication, and Web authentication | The <code>port-security triple-auth-order mac-dot1x-web</code> command was added. |
| Port security: Specifying a domain for port security users redirected to an unavailable URL | The <code>port-security url-unavailable domain</code> command was added. |
| Port security: Clearing port security statistics. | The <code>reset port-security statistics</code> command was added. |
| Object group features | All Object group commands were newly added. |
| Attack detection and prevention: Clearing attack detection and prevention statistics for an interface | The <code>reset attack-defense statistics interface interface-type interface-number</code> command was added. |
| Attack detection and prevention: Displaying information about IPv4 scanning attackers | The <code>interface interface-type interface-number</code> option was added to the <code>display attack-defense scan attacker ip</code> command. |
| Attack detection and prevention: Displaying information about IPv6 scanning attackers | The <code>interface interface-type interface-number</code> option was added to the <code>display attack-defense scan attacker ipv6</code> command. |
| Attack detection and prevention: Displaying information about IPv4 scanning attack victims | The <code>interface interface-type interface-number</code> option was added to the <code>display attack-defense scan victim ip</code> command. |
| Attack detection and prevention: Displaying information about IPv6 scanning attack victims | The <code>interface interface-type interface-number</code> option was added to the <code>display attack-defense scan victim ipv6</code> command. |
| ND attack defense: Ignoring ingress ports of ND packets in ND attack detection | The <code>ipv6 nd detection port-match-ignore</code> command was added. |
| ND attack defense: Displaying ND keepalive | The <code>display ipv6 nd scan keepalive</code> |

| Feature | Command changes |
|---|---|
| entries | The entry [interface <i>interface-type</i> <i>interface-number</i>] [count] command was added. |
| ND attack defense: Displaying statistics about NS packets sent to the IPv6 addresses in offline keepalive entries | The display ipv6 nd scan keepalive statistics [slot <i>slot-number</i>] [interface <i>interface-type</i> <i>interface-number</i>] command was added. |
| ND attack defense: Setting the aging time for ND keepalive entries | The ipv6 nd scan keepalive aging-time <i>time</i> command was added. |
| ND attack defense: Enabling the ND keepalive entry scanning feature | The ipv6 nd scan keepalive enable command was added. |
| ND attack defense: Setting the NS packet sending rate for keepalive entry scanning | The ipv6 nd scan keepalive send-rate <i>pps</i> command was added. |
| ND attack defense: Clearing statistics about NS packets sent to the IPv6 addresses in offline keepalive entries | The reset ipv6 nd scan keepalive statistics [slot <i>slot-number</i>] command was added. |
| ND attack defense: Enabling SNMP notifications for ND | The snmp-agent trap enable nd [entry-limit local-conflict nd-miss rate-limit] * command was added. |
| Password control: Setting the maximum number of blacklist entries for a user account | The password-control per-user blacklist-limit command was added. |
| PKI: Enabling local certificate expiration notification | The pki certificate logging enable command was added. |
| SSH: Displaying records for SSH user login exceptions | The display ssh exception-record command was added. |
| SSH: Setting the maximum number of records for SSH user login exceptions | The ssh exception-record max-number command was added. |
| SSH: Setting alarm and recovery thresholds for SSH user login failures in the specified statistics period | The ssh server login-failed threshold-alarm command was added. |
| MACsec: Enabling inclusion of the SCI in the SecTAG field of MACsec frames on an interface | The macsec include-sci command was added. |
| Microsegmentation: Configuring microsegmentation | All microsegmentation commands were newly added. |
| IP source guard: Displaying local IPv4SG bindings that can be synchronized by routing protocols | The display ip source binding-local command was added. |
| IP source guard: Displaying remote IPv4SG bindings synchronized by routing protocols | The display ip source binding-remote command was added. |
| IP source guard: Displaying statistics about local and remote IPv4SG bindings that routing protocols synchronize | The display ip source binding statistics command was added. |
| IP source guard: Displaying local IPv6SG bindings that can be synchronized by routing protocols | The display ipv6 source binding-local command was added. |
| IP source guard: Displaying remote IPv6SG | The display ipv6 source |

| Feature | Command changes |
|--|---|
| bindings synchronized by routing protocols | binding-remote command was added. |
| IP source guard: Displaying statistics about local and remote IPv6SG bindings that routing protocols synchronize | The display ipv6 source binding statistics command was added. |
| ARP attack protection: Displaying ARP source suppression entries | The display arp source-suppression cache command was added. |
| ARP attack protection: Setting the ARP packet sending rate for automatic ARP scanning | The arp scan auto send-rate command was added. |
| ARP attack protection: Enabling automatic ARP scanning | The arp scan auto enable command was added. |
| ARP attack protection: Setting the ARP request sending rate for keepalive entry scanning | The arp scan keepalive send-rate command was added. |
| ARP attack protection: Setting the aging time for ARP keepalive entries | The arp scan keepalive aging-time command was added. |
| ARP attack protection: Enabling the ARP keepalive entry scanning feature | The arp scan keepalive enable command was added. |
| ARP attack protection: Displaying ARP keepalive entries | The display arp scan keepalive entry command was added. |
| ARP attack protection: Displaying statistics about ARP requests sent to the IP addresses in offline keepalive entries for an interface | The display arp scan keepalive statistics command was added. |
| ARP attack protection: Clearing statistics about ARP requests sent to the IP addresses in offline keepalive entries | The reset arp scan keepalive statistics command was added. |
| ARP attack protection: Enabling ARP gateway protection for a gateway | The arp filter source command was added. |
| ARP attack protection: Enabling SNMP notifications for ARP | The snmp-agent trap enable arp command was added. |
| IP source guard: Enabling IPv6SG | The ipv6 verify source command was added to VLAN view. |
| IP source guard: Displaying IPv4SG bindings | The remote keyword was added to the display ip source binding command. |
| ARP attack protection: Configuring a user validity check rule | The vsi vsi-name option was added to the arp detection rule command. |
| ARP attack protection: Enabling ARP attack detection logging | The threshold threshold-value option was added to the arp detection log enable command. |

New features: High availability features

Table 9 describes the high availability features added in this software version.

For more information about the features, see *High Availability Configuration Guide* in *HPE 5940 Switch Series Configuration Guides-Release 671x*.

For more information about the commands, see *High Availability Command Reference* in *HPE 5940 Switch Series Command References-Release 671x*.

Table 9 High availability features added in Release 6710

| Feature | Command changes |
|---|--|
| VRRP: Enabling the isolation mode for IPv4 VRRP | The vrrp isolate enable command was newly added. |
| VRRP: Setting the delay time for the IPv4 VRRP group to transition from Initialize to Master or Backup state | The vrrp state-transition-delay command was newly added. |
| VRRP: Enabling the isolation mode for IPv6 VRRP | The vrrp ipv6 isolate enable command was newly added. |
| VRRP: Setting the delay time for the IPv6 VRRP group to transition from Initialize to Master or Backup state | The vrrp ipv6 state-transition-delay command was newly added. |
| BFD: Setting the delay timer for BFD to notify upper-layer protocols of session establishment failures | The bfd init-fail-timer command was newly added. |
| BFD: Setting the minimum interval for receiving multihop BFD echo packets | The bfd multi-hop min-echo-receive-interval command was newly added. |
| BFD: Creating a static BFD session and entering its view, or entering the view of an existing static BFD session | The bfd static command was newly added. |
| BFD: Specifying the TTL value for BFD packets | The bfd ttl command was newly added. |
| BFD: Specifying the local and remote discriminators for a static BFD session | The discriminator command was newly added. |
| BFD: Displaying the TTL values for BFD packets | The display bfd ttl command was newly added. |
| Track: Creating a track entry associated with a static BFD session and entering track entry view, or entering the view of an existing track entry | The track bfd static command was newly added. |
| Track: Creating a track entry associated with M-LAG MAD state and entering track entry view, or entering the view of an existing track entry | The track mlag-mad-status command was newly added. |
| Track: Disabling checking whether the echo response receiving interface is consistent with the echo packet output interface | The ignore-rx-interface keyword was added to the track bfd echo command. |

New features: Network management and monitoring features

Table 10 describes the network management and monitoring features added in this software version.

For more information about the features, see *Network Management and Monitoring Configuration Guide* in *HPE 5940 Switch Series Configuration Guides-Release 671x*.

For more information about the commands, see *Network Management and Monitoring Command Reference* in *HPE 5940 Switch Series Command References-Release 671x*.

Table 10 Network management and monitoring features added in Release 6710

| Feature | Command changes |
|---|--|
| NTP: Disallowing control queries from peer devices to the local device | The ntp-service noquery enable command was added. |
| PTP: Displaying brief information about the PTP synchronization path from the GM to the device | The display ptp path-trace command was added. |
| PTP: Displaying historical role change information for PTP ports | The display ptp port-history command was added. |
| PTP: Enable PTP globally | The ptp global enable command was added. |
| PTP: Disabling PTP path tracing | The ptp path-trace disable command was added. |
| sFlow: Configuring sFlow collector information | The dscp dscp-value option was added to the sflow collector command. |
| Performance management | All performance management commands were newly added. |
| System maintenance and debugging: Tracing the path that the packets traverse from source to destination | The -e keyword was added to the following commands: <ul style="list-style-type: none"> tracert tracert ipv6 |
| SNMP: Configuring the target host to which SNMP notifications are sent | The cipher-securityname cipher-security-string option was added to the snmp-agent target-host command. |
| SNMP: Specifying the notification format | The snmp-agent trap format command was added. |
| SNMP: Adding the device serial number (SN) to SNMP notifications sent from the device to the NMS | The snmp-agent trap withsn command was added. |
| EAA: Configuring an automatic email sending action | The following commands were added: <ul style="list-style-type: none"> rtm email domain rtm email max-size rtm email username password action email |
| EAA: Configuring a periodic event | The event period command was added. |
| Process monitoring and maintenance: Displaying free memory block information for a user process | The display process memory fragment free command was added. |
| Process monitoring and maintenance: Displaying used memory block information for a user process | The display process memory fragment used command was added. |
| Information center: Setting the format for logs sent to log hosts | The rfc5424 keyword was added to the info-center format command. |
| Information center: Adding the device serial number to the location field of logs sent to log hosts | The info-center loghost locate-info with-sn command was added. |
| VCF fabric: Specifying the NETCONF | The following commands were added: |

| Feature | Command changes |
|--|--|
| username and password for automated VCF fabric deployment | <ul style="list-style-type: none"> vcf-fabric underlay netconf-username vcf-fabric underlay netconf-password |
| NETCONF: Displaying current NETCONF service status and global NETCONF service statistics | The display netconf service command was added. |
| NETCONF: Displaying NETCONF session status and statistics | The display netconf session command was added. |
| NETCONF: Clearing current global NETCONF service statistics | The reset netconf service statistics command was added. |
| NETCONF: Clearing current NETCONF session statistics | The reset netconf session statistics command was added. |
| NETCONF: Enabling conversion of NETCONF operations to logs and commands | The netconf log xml2cli enable command was added. |
| OpenFlow: Specifying a controller for an OpenFlow switch and configuring the main connection to the controller | The global-ssl keyword was added to the controller address command. |
| Feature | Command changes |
| NTP: Disallowing control queries from peer devices to the local device | The ntp-service noquery enable command was added. |
| PTP: Displaying brief information about the PTP synchronization path from the GM to the device | The display ptp path-trace command was added. |
| PTP: Displaying historical role change information for PTP ports | The display ptp port-history command was added. |
| PTP: Enable PTP globally | The ptp global enable command was added. |
| PTP: Disabling PTP path tracing | The ptp path-trace disable command was added. |
| sFlow: Configuring sFlow collector information | The dscp dscp-value option was added to the sflow collector command. |
| Performance management | All performance management commands were newly added. |
| System maintenance and debugging: Tracing the path that the packets traverse from source to destination | The -e keyword was added to the following commands: <ul style="list-style-type: none"> tracert tracert ipv6 |
| SNMP: Configuring the target host to which SNMP notifications are sent | The cipher-securityname cipher-security-string option was added to the snmp-agent target-host command. |
| SNMP: Specifying the notification format | The snmp-agent trap format command was added. |
| SNMP: Adding the device serial number (SN) to SNMP notifications sent from the device to the NMS | The snmp-agent trap withsn command was added. |
| EAA: Configuring an automatic email sending | The following commands were added: |

| Feature | Command changes |
|---|---|
| action | <ul style="list-style-type: none"> <code>rtm email domain</code> <code>rtm email max-size</code> <code>rtm email username password</code> <code>action email</code> |
| EAA: Configuring a periodic event | The <code>event period</code> command was added. |
| Process monitoring and maintenance: Displaying free memory block information for a user process | The <code>display process memory fragment free</code> command was added. |
| Process monitoring and maintenance: Displaying used memory block information for a user process | The <code>display process memory fragment used</code> command was added. |
| Information center: Setting the format for logs sent to log hosts | The <code>rfc5424</code> keyword was added to the <code>info-center format</code> command. |
| Information center: Adding the device serial number to the location field of logs sent to log hosts | The <code>info-center loghost locate-info with-sn</code> command was added. |
| VCF fabric: Specifying the NETCONF username and password for automated VCF fabric deployment | <p>The following commands were added:</p> <ul style="list-style-type: none"> <code>vcf-fabric underlay netconf-username</code> <code>vcf-fabric underlay netconf-password</code> |
| NETCONF: Displaying current NETCONF service status and global NETCONF service statistics | The <code>display netconf service</code> command was added. |
| NETCONF: Displaying NETCONF session status and statistics | The <code>display netconf session</code> command was added. |
| NETCONF: Clearing current global NETCONF service statistics | The <code>reset netconf service statistics</code> command was added. |
| NETCONF: Clearing current NETCONF session statistics | The <code>reset netconf session statistics</code> command was added. |
| NETCONF: Enabling conversion of NETCONF operations to logs and commands | The <code>netconf log xml2cli enable</code> command was added. |

New features: Telemetry features

Table 11 describes the telemetry features added in this software version.

For more information about the features, see *Telemetry Configuration Guide* in *HPE 5940 Switch Series Configuration Guides-Release 671x*.

For more information about the commands, see *Telemetry Command Reference* in *HPE 5940 Switch Series Command References-Release 671x*.

Table 11 Telemetry features added in Release 6710

| Feature | Command changes |
|--------------------------------|---|
| gRPC: Displaying detailed gRPC | The <code>verbose</code> keyword was added to the <code>display grpc</code> |

| Feature | Command changes |
|---|---|
| information | command. |
| gRPC: Setting the maximum CPU usage for gRPC | The grpc cpu-usage max-percent command was added. |
| gRPC: Displaying sensor paths that have the minimum sampling interval | The display telemetry sensor-path command was added. |
| gRPC: Adding an IPv4 collector to the destination group by its domain name | The domain-name command was added. |
| gRPC: Adding an IPv6 collector to the destination group by its domain name | The ipv6 domain-name command was added. |
| gRPC: Setting the DSCP value for packets sent to collectors | The dscp command was added. |
| gRPC: Enabling TLS to encrypt the gRPC connection between the device and the specified IPv4 collector | The tls keyword was added to the ipv4-address command. |
| gRPC: Enabling TLS to encrypt the gRPC connection between the device and the specified collector | The tls keyword was added to the ipv6-address command. |
| gRPC: Enabling per-row time-stamping for JSON-encoded subscription data | The json row-timestamp enable command was added. |
| gRPC: Setting the data push mode for a subscription | The push-mode command was added. |
| gRPC: Pushing data from specified nodes in a sensor path | The selection-nodes node-list option was added to the sensor path command. |
| gRPC: Specifying that the sensor group collects and pushes data at intervals in milliseconds | The msec keyword was added to the sensor-group command. |
| gRPC: Setting the data push suppression interval for a sensor group | The suppress-time suppress-time option was added to the sensor-group command. |
| gRPC: Creating a gNMI sensor group | The gnmi keyword was added to the sensor-group command. |
| gRPC: Creating a gNMI subscription | The gnmi keyword was added to the subscription command. |

New features: OpenFlow features

Table 12 describes the OpenFlow features added in this software version.

For more information about the features, see *OpenFlow Configuration Guide* in *HPE 5940 Switch Series Configuration Guides-Release 671x*.

For more information about the commands, see *OpenFlow Command Reference* in *HPE 5940 Switch Series Command References-Release 671x*.

Table 12 OpenFlow features added in Release 6710

| Feature | Command changes |
|---|--|
| Specifying the global SSL client policy for TTP controllers | The openflow ttp-controller global-ssl command was added. |

| Feature | Command changes |
|--|---|
| Enabling SNMP notifications for OpenFlow | The snmp-agent trap enable openflow command was added. |
| Enabling the TTP mode for an OpenFlow instance | The ttp enable command was added. |
| Excluding the specified VLANs from the VLANs in which traffic is forwarded in the OpenFlow forwarding process | The openflow normal-forward vlan command was added. |
| OpenFlow: Specifying a controller for an OpenFlow switch and configuring the main connection to the controller | The global-ssl keyword was added to the controller address command. |

New features: VXLAN features

Table 13 describes the VXLAN features added in Release 6710.

For more information about the features, see *VXLAN Configuration Guide* in *HPE 5940 Switch Series Configuration Guides-Release 671x*.

For more information about the commands, see *VXLAN Command Reference* in *HPE 5940 Switch Series Command References-Release 671x*.

Table 13 VXLAN features added in Release 6710

| Feature | Command changes |
|--|---|
| Disabling flooding the ARP requests that do not match any ARP flood suppression entries | The no-broadcast keyword was added to the arp suppression enable command. |
| Setting the ARP flood suppression mode | The arp suppression mode command was added. |
| Displaying information about the multicast groups that contain IGMP host-enabled interfaces for a VPN instance | The vpn-instance vpn-instance-name option was added to the display igmp host group command. |
| Displaying MAC move records for VSIs | The display l2vpn mac-address mac-move command was added. |
| Displaying information about VXLAN tunnel interfaces | The display vxlan tunnel-interface command was added. |
| Disabling flooding the ND requests that do not match any ND flood suppression entries | The no-broadcast keyword was added to the ipv6 nd suppression enable command. |
| Setting the ND flood suppression mode | The ipv6 nd suppression mode command was added. |
| Clearing MAC move records for all VSIs | The reset l2vpn mac-address mac-move command was added. |

New features: EVPN features

Table 14 describes the EVPN features added in Release 6710.

For more information about the features, see *EVPN Configuration Guide* in *HPE 5940 Switch Series Configuration Guides-Release 671x*.

For more information about the commands, see *EVPN Command Reference* in *HPE 5940 Switch Series Command References-Release 671x*.

Table 14 EVPN features added in Release 6710

| Feature | Command changes |
|---|--|
| EVPN VXLAN: Enabling ARP request proxy | The arp proxy-send enable command was added. |
| EVPN VXLAN: Displaying ARP gateway protection configuration for Ethernet service instances | The display arp filter source service-instance [interface interface-type interface-number [service-instance instance-id]] [slot slot-number] command was added. |
| EVPN VXLAN: Displaying BGP EVPN routes | The following parameters were added to the display bgp l2vpn evpn command: <ul style="list-style-type: none"> • s-pmsi • cluster-list • community • ext-community • verbose |
| EVPN VXLAN: Displaying EVPN instance information | The display evpn instance [name instance-name vsi vsi-name] vxlan command was added. |
| EVPN VXLAN: Displaying information about IPv6 peers that are automatically discovered through BGP | The display evpn ipv6 auto-discovery { imet [peer ipv6-address] [vsi vsi-name] mac-ip macip-prefix [nexthop next-hop] [count] } command was added. |
| EVPN VXLAN: Displaying IPv6 EVPN MAC address entries | The display evpn ipv6 route mac [local remote] [vsi vsi-name] [count] command was added. |
| EVPN VXLAN: Configuring ARP mobility event suppression | The display evpn route arp-mobility [public-instance vpn-instance vpn-instance-name] [ip ip-address] command was added. The following parameters were added to the evpn route arp-mobility suppression command: <ul style="list-style-type: none"> • detect-cycle detect-time • detect-threshold move-times • suppression-time • permanent The reset evpn route arp-mobility suppression [public-instance vpn-instance vpn-instance-name [ip ip-address]] command was added. |
| EVPN VXLAN: Configuring MAC mobility event suppression | The following commands were added: <ul style="list-style-type: none"> • display evpn [ipv6] route |

| Feature | Command changes |
|---|---|
| | <pre>mac-mobility [vsi vsi-name] [mac-address mac-address]</pre> <ul style="list-style-type: none"> evpn route mac-mobility suppression [detect-cycle <i>detect-time</i> detect-threshold <i>move-times</i> suppression-time [<i>suppression-time</i> permanent]] * reset evpn route mac-mobility suppression [vsi <i>vsi-name</i> [mac <i>mac-address</i>]] |
| EVPN VXLAN: Configuring ND mobility event suppression | <p>The following commands were added:</p> <ul style="list-style-type: none"> display evpn route nd-mobility [public-instance vpn-instance <i>vpn-instance-name</i>] [ipv6 <i>ipv6-address</i>] evpn route nd-mobility suppression [detect-cycle <i>detect-time</i> detect-threshold <i>move-times</i> suppression-time [<i>suppression-time</i> permanent]] * reset evpn route nd-mobility suppression [public-instance vpn-instance <i>vpn-instance-name</i> [ipv6 <i>ipv6-address</i>]] |
| EVPN VXLAN: Enabling the device to advertise ARP information for the distributed EVPN gateway interfaces through MAC/IP advertisement routes | The evpn mac-ip advertise distributed-gateway command was added. |
| EVPN VXLAN: Enabling local FRR globally for EVPN VXLAN | The evpn multihoming vxlan-frr local command was added. |
| EVPN VXLAN: Setting the static flag for the MAC addresses of centralized gateway interfaces | The evpn route gateway-mac unmovable command was added. |
| EVPN VXLAN: Disabling a VSI interface from learning ARP or ND information that does not belong to its subnet from MAC/IP advertisement routes | The evpn span-segment { arp-learning nd-learning } disable command was added. |
| EVPN VXLAN: Setting the metric of a BGP EVPN route added to a VPN instance's routing table to the metric of the IGP route pointing to the next hop in the original BGP EVPN route | The igp-metric inherit command was added. |
| EVPN VXLAN: Enabling ND request proxy | The ipv6 nd proxy-send enable command was added. |
| EVPN VXLAN: Enabling traffic statistics for the VSIs that are automatically created for L3 VXLAN IDs | The l2vpn statistics vsi l3-vni command was added. |
| EVPN VXLAN: Enabling SMET route reorigination based on the remote VXLAN IDs in VXLAN mappings | The mapping-vni-based smet command was added. |

| Feature | Command changes |
|---|---|
| EVPN VXLAN: Enabling the device to ignore default routes in route recursion | The nexthop recursive-lookup default-route ignore [route-policy route-policy-name] command was added. |
| EVPN VXLAN: Enabling advertising the IP prefix advertisement routes reoriginated for a VPN instance to an IBGP peer or peer group | The peer { group-name ipv4-address [mask-length] } advertise vpn-reoriginate ibgp command was added. |
| EVPN VXLAN: Enabling SNMP notifications for EVPN | The snmp-agent trap enable evpn [mac-mobility-suppression] command was added. |
| MVXLAN features | All MVXLAN commands were newly added. |

New features: Intelligent lossless network features

As from this release, documents about the PFC module move to document set Intelligent Lossless Network.

Table 15 describes the PFC features added in this software version.

For more information about the features, see *Intelligent Lossless Network Configuration Guide in HPE 5940 Switch Series Configuration Guides-Release 671x*.

For more information about the commands, see *Intelligent Lossless Network Command Reference in HPE 5940 Switch Series Command References-Release 671x*.

Table 15 PFC features added in Release 6710

| Feature | Command changes |
|---|--|
| PFC: Configuring the action to take on an interface when the number of PFC deadlock times within the specified detection period exceeds the upper threshold | <ul style="list-style-type: none"> The priority-flow-control deadlock threshold action command was added. The error-down keyword was added to the priority-flow-control deadlock threshold command. |

New features: M-LAG support for DRNI commands

The M-LAG feature was named DRNI in earlier software versions. For compatibility with earlier software versions, the device supports both M-LAG and DRNI commands.

Table 16 Difference between M-LAG and DRNI commands

| Feature name | Keywords | Example |
|--------------|-------------------------------|---|
| M-LAG | m-lag, mlag, peer-link | m-lag system-number system-number As a best practice, use M-LAG commands. |
| DRNI | drni, drmac, ipp | drni system-number system-number To use configuration files created from an earlier software version, use DRNI commands. The system recognizes only the complete syntaxes of DRNI commands. It does not support displaying available |

| | | |
|--|--|---|
| | | keywords and arguments in response to a question mark (?) or automatically completing the last keyword or argument in response to the Tab key. If you execute a DRNI command, the system converts it into the corresponding M-LAG command and saves the M-LAG command in the configuration file. |
|--|--|---|

M-LAG and DRNI commands do not differ in the configuration method or functionality. For more information about the keyword differences, see comparison between M-LAG and DRNI commands see [Modified feature: DRNI term changes](#).

New feature: EVPN instance created in system view

Configuring an EVPN instance created in system view

About this task

You can bind an EVPN instance created in system view to multiple VSIs to simplify configuration.

Restrictions and guidelines

You can bind a VSI only to one EVPN instance.

Procedure

1. Enter system view.
system-view
2. Create an EVPN instance and enter its view.
evpn instance *instance-name*
3. Configure an RD for the EVPN instance.
route-distinguisher *route-distinguisher*
By default, no RD is configured for an EVPN instance.
4. Configure route targets for the EVPN instance.
vpn-target *vpn-target*&<1-8> [**both** | **export-extcommunity** | **import-extcommunity**]
By default, an EVPN instance does not have route targets.
Make sure the following requirements are met:
 - The import targets of the EVPN instance do not match the export targets of the VPN instance associated with the VXLAN or the public instance.
 - The export targets of the EVPN instance do not match the import targets of the VPN instance associated with the VXLAN or the public instance.
5. Return to system view.
quit
6. Enter VSI view.
vsi *vsi-name*
7. Bind the VSI to the EVPN instance.
evpn encapsulation vxlan binding instance *instance-name* **vsi-tag** { *tag-id* | **auto-vxlan** }
By default, a VSI is not bound to an EVPN instance created in system view.

Command reference

New command: `evpn encapsulation vxlan binding instance`

Use `evpn encapsulation vxlan binding instance` to bind a VSI to an EVPN instance.

Use `undo evpn encapsulation vxlan binding instance` to restore the default.

Syntax

```
evpn encapsulation vxlan binding instance instance-name vsi-tag { tag-id  
| auto-vxlan }
```

```
undo evpn encapsulation vxlan binding instance instance-name
```

Views

VSI view

Default

A VSI is not bound to any EVPN instance.

Predefined user roles

network-admin

Parameters

instance-name: Specifies an EVPN instance by its name, a case-sensitive string of 1 to 31 characters.

vsi-tag: Specifies a VSI tag ID. An EVPN instance uses tag IDs to identify VSIs when it is bound to multiple VSIs.

tag-id: Specifies a tag ID in the range of 1 to 16777215.

auto-vxlan: Automatically generates a tag ID based on the VXLAN ID of the VSI.

Usage guidelines

This command is applicable only to EVPN VXLAN.

You can bind a VSI only to one EVPN instance.

This command is mutually exclusive with the `evpn encapsulation` command. You cannot use them together on the same VSI.

Examples

Bind VSI **vpna** to EVPN instance **evpna**.

```
<Sysname> system-view
```

```
[Sysname] vsi vpna
```

```
[Sysname-vsi-vpna] evpn encapsulation vxlan binding instance evpna vsi-tag auto-vxlan
```

Related commands

evpn encapsulation

evpn instance

New command: `evpn instance`

Use `evpn instance` to create an EVPN instance and enter its view, or enter the view of an existing EVPN instance.

Use `undo evpn instance` to delete an EVPN instance.

Syntax

```
evpn instance instance-name
undo evpn instance instance-name
```

Views

System view

Default

No EVPN instance exists.

Predefined user roles

network-admin

Parameters

instance-name: Specifies an EVPN instance name, a case-sensitive string of 1 to 31 characters.

Usage guidelines

You must create an EVPN instance before you can configure EVPN on it.

You can bind an EVPN instance created in system view to multiple VSIs to simplify configuration.

Examples

Create EVPN instance **evpn1** and enter its view.

```
<Sysname> system-view
[Sysname] evpn instance evpn1
[Sysname-evpn-instance-evpn1]
```

Modified command: arp mac-learning disable

Syntax

```
arp mac-learning disable
undo arp mac-learning disable
```

Old views

VSI EVPN instance view

New views

EVPN instance view
VSI EVPN instance view

Change description

The EVPN instance view was added.

Modified command: arp-advertising disable

Syntax

```
arp-advertising disable
undo arp-advertising disable
```

Old views

VSI EVPN instance view

New views

EVPN instance view
VSI EVPN instance view

Change description

The EVPN instance view was added.

Modified command: mac-advertising disable

Syntax

```
mac-advertising disable
undo mac-advertising disable
```

Old views

VSI EVPN instance view

New views

EVPN instance view
VSI EVPN instance view

Change description

The EVPN instance view was added.

Modified command: nd-advertising enable

Syntax

```
nd-advertising enable
undo nd-advertising enable
```

Old views

VSI EVPN instance view

New views

EVPN instance view
VSI EVPN instance view

Change description

The EVPN instance view was added.

Modified command: nd mac-learning disable

Syntax

```
nd mac-learning disable
undo nd mac-learning disable
```

Old views

VSI EVPN instance view

New views

EVPN instance view
VSI EVPN instance view

Change description

The EVPN instance view was added.

Modified command: route-distinguisher

Syntax

```
route-distinguisher route-distinguisher  
undo route-distinguisher
```

Old views

Public instance view

New views

Public instance view

EVPN instance view

Change description

The EVPN instance view was added.

Modified command: vpn-target

Syntax

```
vpn-target vpn-target<1-8> [ both | export-extcommunity |  
import-extcommunity ]  
undo vpn-target { all | vpn-target<1-8> [ both | export-extcommunity |  
import-extcommunity ] }
```

Old views

VPN instance EVPN view

Public instance view

Public instance EVPN view

Public instance IPv4 address family view

Public instance IPv6 address family view

New views

VPN instance EVPN view

Public instance view

Public instance EVPN view

Public instance IPv4 address family view

Public instance IPv6 address family view

EVPN instance view

Change description

The EVPN instance view was added.

Modified feature: IPv6 routes with prefixes longer than 64 bits

Feature change description

Before modification: By default, The device does not support for IPv6 routes with prefixes longer than 64 bits.

After modification: By default, The device supports for IPv6 routes with prefixes longer than 64 bits.

Command changes

Modified command: hardware-resource routing-mode

Syntax

```
hardware-resource routing-mode { ipv6-64 | ipv6-128 }
```

.Views

System view

Change description

Before modification: By default, The device does not support for IPv6 routes with prefixes longer than 64 bits.

After modification: By default, The device supports for IPv6 routes with prefixes longer than 64 bits.

Modified feature: Match criteria in a traffic class

Feature change description

As from this version, you can configure a match criterion to match the input interface of packets in a traffic class.

Modified command: if-match

Syntax

```
if-match match-criteria
```

```
undo if-match match-criteria
```

Views

Traffic class view

Change description

Before modification: You cannot configure a match criterion to match the input interface of packets.

After modification: You can configure a match criterion to match the input interface of packets by specifying the **inbound-interface** *interface-type interface-number* option for the *match-criteria* argument.

Table 2 Available match criteria

| Option | Description |
|--|--|
| inbound-interface <i>interface-type</i> <i>interface-number</i> | <p>Matches an input interface specified by its type and number.</p> <p>If this option is configured in a traffic class with logic AND operator, the traffic class is no longer in effect after the subcard where the input interface resides is removed. After the removed subcard is reinserted, the traffic class takes effect again. If you do not reinsert the subcard, you must delete this traffic class and configure it again as needed. Otherwise, even if you add other match criteria to the traffic class, the traffic class does not take effect again.</p> |

Modified feature: Associating a traffic behavior with a traffic class

Feature change description

As from this version, you can associate a traffic behavior with a traffic class in loose mode. The loose mode takes effect only when you apply a QoS policy to the control plane.

The loose mode can match those packets that cannot be identified, such as ARP packets. Use the loose mode with caution, because it might mistakenly match protocol packets.

Follow these steps to use the loose mode to rate limit the packets of the specified protocol received on the specified interface and sent to the control plane:

1. Create an ACL, and create a rule that matches packets of the specified protocol in the ACL.
2. Create a traffic class with logic AND operator, and configure the **if-match acl** and **if-match inbound-interface** criteria in the traffic class.
3. Create a traffic behavior and configure actions in it.
4. Create a QoS policy, and specify the **mode loose** keyword when you associate the traffic behavior with the traffic class configured above.
5. Apply the QoS to the control plane.

Protocol packets of the specified protocol received on any other interfaces are still rate limited by the protocol packet rate limiting feature when they are sent to the control plane.

Command changes

Modified command: classifier behavior

Old syntax

```
classifier classifier-name behavior behavior-name [ mode { dcbx |
qppb-manipulation } | insert-before before-classifier-name ]
```

New syntax

```
classifier classifier-name behavior behavior-name [ mode { dcbx | loose |
qppb-manipulation } | insert-before before-classifier-name ]
```

Views

QoS policy view

Change description

Before modification: You cannot associate a traffic behavior with a traffic class in loose mode

After modification: You can associate a traffic behavior with a traffic class in loose mode

Modified feature: Displaying the running configuration

Feature change description

As from this release, the running configuration can be displayed by section.

Command changes

Modified command: display current-configuration

Old syntax

```
display current-configuration [ [ configuration [ module-name ] |  
exclude-provision | interface [ interface-type [ interface-number ] ] ] [ all ]  
| slot slot-number ]
```

New syntax

```
display current-configuration [ [ configuration [ module-name ] |  
exclude-provision | interface [ interface-type [ interface-number ] ] ] [ all ]  
| slot slot-number ] [ by-section { begin | exclude | include }  
regular-expression ]
```

Views

Any view

Change description

The following parameters are added to this command:

by-section: Displays the running configuration by section. Each section contains the configuration in a command view and two adjacent sections are separated by a pound sign (#). If you do not specify this keyword, the command displays the running configuration by line.

begin: Displays the first configuration section that matches the specified regular expression and all subsequent configuration sections.

exclude: Displays all configuration sections not matching the specified regular expression.

include: Displays all configuration sections matching the specified regular expression.

regular-expression: Specifies a regular expression to filter the configuration sections, a case-sensitive string of 1 to 256 characters. A section matches the specified regular expression if it contains command lines that match the specified regular expression.

Modified feature: Displaying the contents of the configuration file for the next system startup

Feature change description

As from this release, the contents of the configuration file for the next system startup can be displayed by section.

Command changes

Modified command: display saved-configuration

Old syntax

```
display saved-configuration
```

New syntax

```
display saved-configuration [ by-section { begin | exclude | include }  
regular-expression ]
```

Views

Any view

Change description

The following parameters are added to this command after modification:

by-section: Displays the configuration used at the next startup by section. Each section contains the configuration in a command view and two adjacent sections are separated by a pound sign (#). If you do not specify this keyword, the command displays the configuration used at the next startup by line.

begin: Displays the first configuration section that matches the specified regular expression and all subsequent configuration sections.

exclude: Displays all configuration sections not matching the specified regular expression.

include: Displays all configuration sections matching the specified regular expression.

regular-expression: Specifies a regular expression to filter the configuration sections, a case-sensitive string of 1 to 256 characters. A section matches the specified regular expression if it contains command lines that match the specified regular expression.

Modified feature: Optimized display of BGP BMP server information

Feature change description

As from this release, the output of the **display bgp bmp server** command includes the following information:

- Interval at which BGP sends statistics information to the BMP server.
- Type of the routes that BGP reports to the BMP server.
- Whether the BMP client sends peer up notifications with mode flags to the BMP server.

- Whether the BMP client sends peer down notifications with mode flags to the BMP server.

Command changes

Modified command: display bgp bmp server

Syntax

```
display bgp [ instance instance-name ] bmp server server-number
```

Views

Any view

Change description

After modification, the following fields are added to the output of this command:

- **Statistics report interval:** Interval (in seconds) at which BGP sends statistics information to the BMP server.
- **Reported route mode:** Type of routes that BGP sends to the BMP server:
 - **adj-rib-in**—Send routes received from the monitored peer or peer group to the BMP server.
 - **adj-rib-out**—Send routes advertised to the monitored peer or peer group to the BMP server.
 - **pre-policy**—Send routes to the BMP server without route filtering.
 - **post-policy**—Send routes to the BMP server after route filtering.
 - **both**—Send both filtered and unfiltered routes to the BMP server.
 - **loc-rib**—Send the optimal routes in the routing table to the BMP server.
- **Pu-monitor-mode:** Whether the peer up notifications that the BMP client sends to the BMP server carry the mode flag.
 - **Enabled**—Carry the mode flag.
 - **Disabled**—Do not carry the mode flag.
- **Pd-monitor-mode:** Whether the peer down notifications that the BMP client sends to the BMP server carry the mode flag.
 - **Enabled**—Carry the mode flag.
 - **Disabled**—Do not carry the mode flag.

Examples

```
# Display information about BMP server 1.
<Sysname> display bgp bmp server 1
BMP server number: 1
Server VPN instance name: vpna
Server address: 100.1.1.1  Server port: 6895
Client address: 100.1.1.2  Client port: 21452
BMP server state: Connected  Up for 00h41m53s
TCP source interface has been configured
Statistics report interval: 5s
Reported route mode: adj-rib-in pre-policy
Pu-monitor-mode: Enabled
Pd-monitor-mode: Enabled
Message statistics:
Total messages sent: 15
      INITIATION: 1
      TERMINATION: 0
```

```

STATS-REPORT: 0
    PEER-UP: 4
    PEER-DOWN: 3
    ROUTE-MON: 7
BGP peers monitored by BMP server:
10.1.1.1

```

Modified feature: Disabling BGP session establishment with peers and peer groups

Feature change description

As from this release, the value range for the **graceful** *graceful-time* option in the **ignore all-peers** command and the **peer ignore** command is changed. The new value range is 0 to 65535. If you set the value to 0 for the *graceful-time* argument, the device sends low-priority routes to peers or peer groups and does not tear down BGP sessions to the peers or peer groups.

Command changes

Modified command: ignore all-peers

Syntax

```

ignore all-peers [ graceful graceful-time { community { community-number
| aa:nn } | local-preference preference | med med } * ]

```

Views

BGP instance view

Change description

- Before modification: The value range for the *graceful-time* argument is 60 to 65535 seconds.
- After modification: The value range for the *graceful-time* argument is 0 to 65535 seconds. If you set the value to 0 for this argument, the device does not tear down BGP sessions to peers and peer groups.

Modified command: interface-peer/peer ignore

Syntax

```

peer { group-name | ipv4-address [ mask-length ] | ipv6-address
[ prefix-length ] | link-local-address interface interface-type
interface-number } ignore [ graceful graceful-time { community
{ community-number | aa:nn } | local-preference preference | med med } * ]

interface-peer interface-type interface-number ignore [ graceful
graceful-time { community { community-number | aa:nn } | local-preference
preference | med med } * ]

```

Views

BGP instance view

BGP-VPN instance view

Change description

- Before modification: The value range for the *graceful-time* argument is 60 to 65535 seconds.
- After modification: The value range for the *graceful-time* argument is 0 to 65535 seconds. If you set the value to 0 for this argument, the device does not tear down BGP sessions to the specified peers or peer group.

Modified feature: Optimizations to VXLAN command output

Feature change description

In this software version, the following optimizations were made to VXLAN display commands:

- The **display l2vpn vsi** command displays both input and output packet rates for a VSI.
- The **display vxlan tunnel** command displays the total number of VXLAN tunnels for a VXLAN, the source and destination addresses of each VXLAN tunnel, and the outgoing VXLAN ID.

Command changes

Modified command: display l2vpn vsi

Syntax

```
display l2vpn vsi [ name vsi-name ] [ verbose ]
```

Views

Any view

Change description

The following fields were added to the command output:

- **Input Rate:** Incoming traffic rate for a VSI, in bps and pps.
- **Output Rate:** Outgoing traffic rate for a VSI, in bps and pps.

Modified command: display vxlan tunnel

Syntax

```
display vxlan tunnel [ vxlan-id vxlan-id ]
```

Views

Any view

Change description

The following fields were added to the command output:

- **Total number of VXLAN tunnels:** The total number of VXLAN tunnels assigned to a VXLAN.
- **Source:** Tunnel source address.
- **Destination:** Tunnel destination address.
- **Out VNI:** Remote VXLAN ID mapped to the local VXLAN ID by using the **mapping vni** command. If no remote VXLAN ID is configured, this field displays a hyphen (-).

Modified feature: Sharing of VSI interfaces among VSIs

Before modification: Multiple VSIs can share a VSI interface.

After modification: A VSI interface can be assigned to only one VSI.

Modified feature: Enabling L2TP for the specified protocol

Feature change description

As from this version, a Layer 2 Ethernet interface supports the **user-defined** keyword and a Layer 2 aggregate interface supports the **cdp**, **lACP**, **lldp**, **pagp**, **udld**, and **user-defined protocol-name** parameters when you enable L2TP for the specified protocol.

Command changes

Modified command: l2protocol tunnel dot1q

Old syntax

In Layer 2 Ethernet interface view:

```
l2protocol { cdp | dldp | eoam | gvrp | lACP | lldp | mvrp | pagp | pvst | stp |  
udld | vtp } tunnel dot1q
```

In Layer 2 aggregate interface view:

```
l2protocol { gvrp | mvrp | pvst | stp | vtp } tunnel dot1q
```

New syntax

In Layer 2 Ethernet interface view:

```
l2protocol { cdp | dldp | eoam | gvrp | lACP | lldp | mvrp | pagp | pvst | stp |  
udld | user-defined protocol-name | vtp } tunnel dot1q
```

In Layer 2 aggregate interface view:

```
l2protocol { cdp | gvrp | lACP | lldp | mvrp | pagp | pvst | stp | udld |  
user-defined protocol-name | vtp } tunnel dot1q
```

Change description

Before modification: A Layer 2 Ethernet interface does not support the **user-defined** keyword. A Layer 2 aggregate interface does not support the **cdp**, **lACP**, **lldp**, **pagp**, **udld**, or **user-defined protocol-name** parameters.

After modification: A Layer 2 Ethernet interface supports the **user-defined** keyword. A Layer 2 aggregate interface supports the **cdp**, **lACP**, **lldp**, **pagp**, **udld**, and **user-defined protocol-name** parameters.

cdp: Specifies CDP.

lACP: Specifies LACP.

lldp: Specifies LLDP.

pagp: Specifies PAGP.

udld: Specifies UDLD.

user-defined: Specifies a user-defined Layer 2 protocol.

Modified feature: Creating a service loopback group and specify its service type

Feature change description

From this release, you can specify the service loopback group type as **inter-vpn-fwd** (inter-VPN forwarding service) when you create a service loopback group.

Command changes

Modified command: service-loopback group

Old syntax

```
service-loopback group group-id type { { multicast-tunnel | tunnel } * |  
multiport | vsi-gateway }
```

New syntax

```
service-loopback group group-id type { inter-vpn-fwd | { multicast-tunnel  
| tunnel } * | multiport | vsi-gateway }
```

Views

System view

Change description

The **inter-vpn-fwd** keyword was added. This keyword specifies the inter-VPN forwarding service for the service loopback group to be created.

Modified feature: Displaying LSP information

Feature change description

As from this release, the device supports displays SRLSPs established by using SR-TE.

Command changes

Modified command: display mpls lsp

Old syntax

```
display mpls lsp [ egress | in-label label-value | ingress |  
outgoing-interface interface-type interface-number | protocol { bgp |  
isis | ldp | local | rsvp-te | static | static-cr } | transit ] [ vpn-instance  
vpn-instance-name ] [ ipv4-address mask-length | ipv6 [ ipv6-address  
prefix-length ] ] [ verbose ]
```

New syntax

```
display mpls lsp [ egress | in-label label-value | ingress |  
outgoing-interface interface-type interface-number | protocol { bgp |
```

```
isis | ldp | local | sr-te | rsvp-te | static | static-cr } | transit ]  
[ vpn-instance vpn-instance-name ] [ ipv4-address mask-length | ipv6  
[ ipv6-address prefix-length ] ] [ verbose ]
```

Views

Any view

Change description

The **sr-te** keyword was added.

sr-te: Displays SRLSPs established by using SR-TE.

Modified feature: Configuring the signaling protocol for an MPLS TE tunnel interface

Feature change description

As from this release, the device supports specifying Segment Routing as the signaling protocol for establishing an MPLS TE tunnel.

Command changes

Modified command: mpls te signaling

Old syntax

```
mpls te signaling { rsvp-te | static }
```

New syntax

```
mpls te signaling { rsvp-te | static | segment-routing }
```

Views

Tunnel interface view

Change description

The **segment-routing** keyword was added.

segment-routing: Uses the Segment Routing protocol to establish a tunnel.

Modified feature: Configuring the tunnel selection order and set the number of tunnels for load balancing

Feature change description

As from this release, a tunnel policy can select SRLSP tunnels and supports using only the same type of tunnels for load balancing.

Command changes

Modified command: mpls te signaling

Old syntax

```
select-seq { cr-lsp | lsp } * load-balance-number number
```

New syntax

```
select-seq [ strict ] { cr-lsp | lsp | sr-lsp } * load-balance-number number
```

Views

Tunnel policy view

Change description

The **strict** and **sr-lsp** keywords were added.

strict: Uses the same type of tunnels for load balancing. If you do not specify this keyword, the device can use a hybrid of the specified types of tunnels for load balancing.

sr-lsp: Uses SRLSP tunnels.

Modified feature: Verifying MPLS LSP connectivity for an IPv4 prefix

Feature change description

As from this version, the **ping mpls ipv4** command supports the **fec-type**, **generic**, **isis**, and **ldp** keywords.

Command changes

Modified command: ping mpls ipv4

Old syntax

```
ping mpls [ -a source-ip | -c count | -exp exp-value | -h ttl-value | -m wait-time | -r reply-mode | -rtos tos-value | -s packet-size | -t time-out | -v ] * ipv4 ipv4-address mask-length [ destination start-address [ end-address [ address-increment ] ] ]
```

New syntax

```
ping mpls [ -a source-ip | -c count | -exp exp-value | -h ttl-value | -m wait-time | -r reply-mode | -rtos tos-value | -s packet-size | -t time-out | -v ] * ipv4 ipv4-address mask-length [ destination start-address [ end-address [ address-increment ] ] ] [ fec-type { generic | isis | ldp } ]
```

Views

Any view

Change description

The **fec-type**, **generic**, **isis**, and **ldp** keywords were added to the command.

fec-type: Specifies the FEC type of the MPLS LSPs to be verified. If you do not specify this keyword, the FEC type is the protocol type of the specified FEC in the routing table.

generic: Represents LSPs of any FEC type.

is-is: Represents IS-IS signaled SRLSPs.

ldp: Represents LDP signaled LSPs.

Modified feature: Tracing MPLS LSPs from the ingress node to the egress node for an IPv4 prefix

Feature change description

As from this version, the **tracert mpls ipv4** command supports the **fec-type**, **generic**, **isis**, and **ldp** keywords.

Command changes

Modified command: tracert mpls ipv4

Old syntax

```
tracert mpls [ -a source-ip | -exp exp-value | -h ttl-value | -r reply-mode  
| -rtos tos-value | -t time-out | -v | fec-check ] * ipv4 ipv4-address  
mask-length [ destination start-address [ end-address  
[ address-increment ] ] ]
```

New syntax

```
tracert mpls [ -a source-ip | -exp exp-value | -h ttl-value | -r reply-mode  
| -rtos tos-value | -t time-out | -v | fec-check ] * ipv4 ipv4-address  
mask-length [ destination start-address [ end-address  
[ address-increment ] ] ] [ fec-type { generic | isis | ldp } ]
```

Views

Any view

Change description

The **fec-type**, **generic**, **isis**, and **ldp** keywords were added to the command.

fec-type: Specifies the FEC type to be carried in an MPLS echo request packet. If you do not specify this keyword, the FEC type is the protocol type of the specified FEC in the routing table.

generic: Represents LSPs of any FEC type.

is-is: Represents IS-IS signaled SRLSPs.

ldp: Represents LDP signaled LSPs.

Modified feature: Creating a local site

Feature change description

As from this version, the value range for the local site ID changes from 0-65535 to 0-256.

Command changes

Modified command: `site`

Syntax

```
site site-id [ range range-value ] [ default-offset default-offset ]  
undo site site-id
```

Views

Auto-discovery VSI BGP signaling view

Change description

Before modification: The value range for the *site-id* argument is 0 to 65535.

After modification: The value range for the *site-id* argument is 0 to 256

Modified feature: Enabling link flapping protection on an interface

Feature change description

As from this version, the value ranges for the level-1 link flapping detection interval and level-1 link flapping detection threshold change, and you can configure the level-2 link flapping detection interval and level-2 link flapping detection threshold.

Command changes

Modified command: `port link-flap protect enable`

Old syntax

```
port link-flap protect enable [ interval interval | threshold threshold ] *
```

New syntax

```
port link-flap protect enable [ interval interval | threshold threshold ] *  
[ second-interval second-interval second-threshold second-threshold ]
```

Views

Ethernet interface view

Change description

Before modification: The value range for the level-1 link flapping detection interval is 10 to 60 seconds. The value range for the level-1 link flapping detection threshold is 5 to 10. You cannot configure the level-2 link flapping detection interval or level-2 link flapping detection threshold.

After modification: The value range for the level-1 link flapping detection interval is 5 to 86400 seconds. The value range for the level-1 link flapping detection threshold is 2 to 1200. You can configure the level-2 link flapping detection interval and level-2 link flapping detection threshold.

Modified feature: AAA methods in an ISP domain

Feature change description

As from this release, remote AAA methods are supported in an ISP domain:

- For default user accounting, RADIUS and HWTACACS are supported.
- For LAN user accounting, RADIUS is supported.
- For login user accounting, RADIUS and HWTACACS are supported.
- For portal user accounting, RADIUS is supported.
- For default user authentication, RADIUS, HWTACACS, and LDAP are supported.
- For LAN user authentication, RADIUS and LDAP are supported.
- For login user authentication, RADIUS, HWTACACS, and LDAP are supported.
- For portal user authentication, RADIUS and LDAP are supported.
- For default user authorization, RADIUS and HWTACACS are supported.
- For LAN user authorization, RADIUS is supported.
- For login user authorization, RADIUS and HWTACACS are supported.
- For portal user authorization, RADIUS is supported.

Command reference

Modified command: accounting default

Old syntax

In non-FIPS mode:

```
accounting default { hwtacacs-scheme hwtacacs-scheme-name [ radius-scheme radius-scheme-name ] [ local ] [ none ] | local [ none ] | none | radius-scheme radius-scheme-name [ hwtacacs-scheme hwtacacs-scheme-name ] [ local ] [ none ] }
```

```
undo accounting default
```

In FIPS mode:

```
accounting default { hwtacacs-scheme hwtacacs-scheme-name [ radius-scheme radius-scheme-name ] [ local ] | local | radius-scheme radius-scheme-name [ hwtacacs-scheme hwtacacs-scheme-name ] [ local ] }
```

```
undo accounting default
```

New syntax

In non-FIPS mode:

```
accounting default { hwtacacs-scheme hwtacacs-scheme-name [ radius-scheme radius-scheme-name ] [ local ] [ none ] | local [ radius-scheme radius-scheme-name | hwtacacs-scheme hwtacacs-scheme-name ] * [ none ] | none | radius-scheme radius-scheme-name [ hwtacacs-scheme hwtacacs-scheme-name ] [ local ] [ none ] }
```

```
undo accounting default
```

In FIPS mode:

```

accounting default { hwtacacs-scheme hwtacacs-scheme-name [ radius-scheme
radius-scheme-name ] [ local ] | local [ radius-scheme radius-scheme-name
| hwtacacs-scheme hwtacacs-scheme-name ] * | radius-scheme
radius-scheme-name [ hwtacacs-scheme hwtacacs-scheme-name ] [ local ] }
undo accounting default

```

Views

ISP domain view

Modified command: accounting lan-access

Old syntax

In non-FIPS mode:

```

accounting lan-access { broadcast radius-scheme radius-scheme-name1
radius-scheme radius-scheme-name2 [ local ] [ none ] | local [ none ] | none
| radius-scheme radius-scheme-name [ local ] [ none ] }
undo accounting lan-access

```

In FIPS mode:

```

accounting lan-access { broadcast radius-scheme radius-scheme-name1
radius-scheme radius-scheme-name2 [ local ] | local | radius-scheme
radius-scheme-name [ local ] }
undo accounting lan-access

```

New syntax

In non-FIPS mode:

```

accounting lan-access { broadcast radius-scheme radius-scheme-name1
radius-scheme radius-scheme-name2 [ local ] [ none ] | local
[ radius-scheme radius-scheme-name ] [ none ] | none | radius-scheme
radius-scheme-name [ local ] [ none ] }
undo accounting lan-access

```

In FIPS mode:

```

accounting lan-access { broadcast radius-scheme radius-scheme-name1
radius-scheme radius-scheme-name2 [ local ] | local [ radius-scheme
radius-scheme-name ] | radius-scheme radius-scheme-name [ local ] }
undo accounting lan-access

```

Views

ISP domain view

Modified command: accounting login

Old syntax

In non-FIPS mode:

```

accounting login { hwtacacs-scheme hwtacacs-scheme-name [ radius-scheme
radius-scheme-name ] [ local ] [ none ] | local [ none ] | none |
radius-scheme radius-scheme-name [ hwtacacs-scheme hwtacacs-scheme-name ]
[ local ] [ none ] }
undo accounting login

```

In FIPS mode:

```
accounting login { hwtacacs-scheme hwtacacs-scheme-name [ radius-scheme
radius-scheme-name ] [ local ] | local | radius-scheme radius-scheme-name
[ hwtacacs-scheme hwtacacs-scheme-name ] [ local ] }

undo accounting login
```

New syntax

In non-FIPS mode:

```
accounting login { hwtacacs-scheme hwtacacs-scheme-name [ radius-scheme
radius-scheme-name ] [ local ] [ none ] | local [ radius-scheme
radius-scheme-name | hwtacacs-scheme hwtacacs-scheme-name ] * [ none ]
| none | radius-scheme radius-scheme-name [ hwtacacs-scheme
hwtacacs-scheme-name ] [ local ] [ none ] }

undo accounting login
```

In FIPS mode:

```
accounting login { hwtacacs-scheme hwtacacs-scheme-name [ radius-scheme
radius-scheme-name ] [ local ] | local [ radius-scheme radius-scheme-name
| hwtacacs-scheme hwtacacs-scheme-name ] * | radius-scheme
radius-scheme-name [ hwtacacs-scheme hwtacacs-scheme-name ] [ local ] }

undo accounting login
```

Views

ISP domain view

Modified command: accounting portal

Old syntax

In non-FIPS mode:

```
accounting portal { broadcast radius-scheme radius-scheme-name1
radius-scheme radius-scheme-name2 [ local ] [ none ] | local [ none ] | none
| radius-scheme radius-scheme-name [ local ] [ none ] }

undo accounting portal
```

In FIPS mode:

```
accounting portal { broadcast radius-scheme radius-scheme-name1
radius-scheme radius-scheme-name2 [ local ] | local | radius-scheme
radius-scheme-name [ local ] }

undo accounting portal
```

New syntax

In non-FIPS mode:

```
accounting portal { broadcast radius-scheme radius-scheme-name1
radius-scheme radius-scheme-name2 [ local ] [ none ] | local
[ radius-scheme radius-scheme-name ] [ none ] | none | radius-scheme
radius-scheme-name [ local ] [ none ] }

undo accounting portal
```

In FIPS mode:

```

accounting portal { broadcast radius-scheme radius-scheme-name1
radius-scheme radius-scheme-name2 [ local ] | local [ radius-scheme
radius-scheme-name ] | radius-scheme radius-scheme-name [ local ] }

undo accounting portal

```

Views

ISP domain view

Modified command: authentication default

Old syntax

In non-FIPS mode:

```

authentication default { hwtacacs-scheme hwtacacs-scheme-name
[ radius-scheme radius-scheme-name ] [ local ] [ none ] | ldap-scheme
ldap-scheme-name [ local ] [ none ] | local [ none ] | none | radius-scheme
radius-scheme-name [ hwtacacs-scheme hwtacacs-scheme-name ] [ local ]
[ none ] }

undo authentication default

```

In FIPS mode:

```

authentication default { hwtacacs-scheme hwtacacs-scheme-name
[ radius-scheme radius-scheme-name ] [ local ] | ldap-scheme
ldap-scheme-name [ local ] | local | radius-scheme radius-scheme-name
[ hwtacacs-scheme hwtacacs-scheme-name ] [ local ] }

undo authentication default

```

New syntax

In non-FIPS mode:

```

authentication default { hwtacacs-scheme hwtacacs-scheme-name
[ radius-scheme radius-scheme-name ] [ local ] [ none ] | ldap-scheme
ldap-scheme-name [ local ] [ none ] | local [ radius-scheme
radius-scheme-name | hwtacacs-scheme hwtacacs-scheme-name ] * [ none ] |
local [ ldap-scheme ldap-scheme-name ] [ none ] | none | radius-scheme
radius-scheme-name [ hwtacacs-scheme hwtacacs-scheme-name ] [ local ]
[ none ] }

undo authentication default

```

In FIPS mode:

```

authentication default { hwtacacs-scheme hwtacacs-scheme-name
[ radius-scheme radius-scheme-name ] [ local ] | ldap-scheme
ldap-scheme-name [ local ] | local [ radius-scheme radius-scheme-name |
hwtacacs-scheme hwtacacs-scheme-name ] * | local [ ldap-scheme
ldap-scheme-name ] | radius-scheme radius-scheme-name [ hwtacacs-scheme
hwtacacs-scheme-name ] [ local ] }

undo authentication default

```

Views

ISP domain view

Modified command: authentication lan-access

Old syntax

In non-FIPS mode:

```
authentication lan-access { ldap-scheme ldap-scheme-name [ local ] [ none ]  
| local [ none ] | none | radius-scheme radius-scheme-name [ local ] [ none ] }  
  
undo authentication lan-access
```

In FIPS mode:

```
authentication lan-access { ldap-scheme ldap-scheme-name [ local ] | local  
| radius-scheme radius-scheme-name [ local ] }  
  
undo authentication lan-access
```

New syntax

In non-FIPS mode:

```
authentication lan-access { ldap-scheme ldap-scheme-name [ local ] [ none ]  
| local [ ldap-scheme ldap-scheme-name | radius-scheme  
radius-scheme-name ] [ none ] | none | radius-scheme radius-scheme-name  
[ local ] [ none ] }  
  
undo authentication lan-access
```

In FIPS mode:

```
authentication lan-access { ldap-scheme ldap-scheme-name [ local ] | local  
[ ldap-scheme ldap-scheme-name | radius-scheme radius-scheme-name ] |  
radius-scheme radius-scheme-name [ local ] }  
  
undo authentication lan-access
```

Views

ISP domain view

Modified command: authentication login

Old syntax

In non-FIPS mode:

```
authentication login { hwtacacs-scheme hwtacacs-scheme-name  
[ radius-scheme radius-scheme-name ] [ local ] [ none ] | ldap-scheme  
ldap-scheme-name [ local ] [ none ] | local | none | radius-scheme  
radius-scheme-name [ hwtacacs-scheme hwtacacs-scheme-name ] [ local ]  
[ none ] }  
  
undo authentication login
```

In FIPS mode:

```
authentication login { hwtacacs-scheme hwtacacs-scheme-name  
[ radius-scheme radius-scheme-name ] [ local ] | ldap-scheme  
ldap-scheme-name [ local ] | local | radius-scheme radius-scheme-name  
[ hwtacacs-scheme hwtacacs-scheme-name ] [ local ] }  
  
undo authentication login
```

New syntax

In non-FIPS mode:

```

authentication login { hwtacacs-scheme hwtacacs-scheme-name
[ radius-scheme radius-scheme-name ] [ local ] [ none ] | ldap-scheme
ldap-scheme-name [ local ] [ none ] | local [ radius-scheme
radius-scheme-name | hwtacacs-scheme hwtacacs-scheme-name ] * [ none ] |
local [ ldap-scheme ldap-scheme-name ] [ none ] | none | radius-scheme
radius-scheme-name [ hwtacacs-scheme hwtacacs-scheme-name ] [ local ]
[ none ] }

```

```
undo authentication login
```

In FIPS mode:

```

authentication login { hwtacacs-scheme hwtacacs-scheme-name
[ radius-scheme radius-scheme-name ] [ local ] | ldap-scheme
ldap-scheme-name [ local ] | local [ radius-scheme radius-scheme-name |
hwtacacs-scheme hwtacacs-scheme-name ] * | local [ ldap-scheme
ldap-scheme-name ] | radius-scheme radius-scheme-name [ hwtacacs-scheme
hwtacacs-scheme-name ] [ local ] }

```

```
undo authentication login
```

Views

ISP domain view

Modified command: authentication portal

Old syntax

In non-FIPS mode:

```

authentication portal { ldap-scheme ldap-scheme-name [ local ] [ none ] |
local [ none ] | none | radius-scheme radius-scheme-name [ local ] [ none ] }

```

```
undo authentication portal
```

In FIPS mode:

```

authentication portal { ldap-scheme ldap-scheme-name [ local ] | local |
radius-scheme radius-scheme-name [ local ] }

```

```
undo authentication portal
```

New syntax

In non-FIPS mode:

```

authentication portal { ldap-scheme ldap-scheme-name [ local ] [ none ] |
local [ ldap-scheme ldap-scheme-name | radius-scheme radius-scheme-name ]
[ none ] | none | radius-scheme radius-scheme-name [ local ] [ none ] }

```

```
undo authentication portal
```

In FIPS mode:

```

authentication portal { ldap-scheme ldap-scheme-name [ local ] | local
[ ldap-scheme ldap-scheme-name | radius-scheme radius-scheme-name ] |
radius-scheme radius-scheme-name [ local ] }

```

```
undo authentication portal
```

Views

ISP domain view

Modified command: authorization default

Old syntax

In non-FIPS mode:

```
authorization default { hwtacacs-scheme hwtacacs-scheme-name
[ radius-scheme radius-scheme-name ] [ local ] [ none ] | local [ none ] | none
| radius-scheme radius-scheme-name [ hwtacacs-scheme
hwtacacs-scheme-name ] [ local ] [ none ] }
```

```
undo authorization default
```

In FIPS mode:

```
authorization default { hwtacacs-scheme hwtacacs-scheme-name
[ radius-scheme radius-scheme-name ] [ local ] | local | radius-scheme
radius-scheme-name [ hwtacacs-scheme hwtacacs-scheme-name ] [ local ] }
```

```
undo authorization default
```

New syntax

In non-FIPS mode:

```
authorization default { hwtacacs-scheme hwtacacs-scheme-name
[ radius-scheme radius-scheme-name ] [ local ] [ none ] | local
[ radius-scheme radius-scheme-name | hwtacacs-scheme
hwtacacs-scheme-name ] * [ none ] | none | radius-scheme radius-scheme-name
[ hwtacacs-scheme hwtacacs-scheme-name ] [ local ] [ none ] }
```

```
undo authorization default
```

In FIPS mode:

```
authorization default { hwtacacs-scheme hwtacacs-scheme-name
[ radius-scheme radius-scheme-name ] [ local ] | local [ radius-scheme
radius-scheme-name | hwtacacs-scheme hwtacacs-scheme-name ] * |
radius-scheme radius-scheme-name [ hwtacacs-scheme hwtacacs-scheme-name ]
[ local ] }
```

```
undo authorization default
```

Views

ISP domain view

Modified command: authorization lan-access

Old syntax

In non-FIPS mode:

```
authorization lan-access { local [ none ] | none | radius-scheme
radius-scheme-name [ local ] [ none ] }
```

```
undo authorization lan-access
```

In FIPS mode:

```
authorization lan-access { local | radius-scheme radius-scheme-name
[ local ] }
```

```
undo authorization lan-access
```

New syntax

In non-FIPS mode:

```

authorization lan-access { local [ radius-scheme radius-scheme-name ]
[ none ] | none | radius-scheme radius-scheme-name [ local ] [ none ] }
undo authorization lan-access

```

In FIPS mode:

```

authorization lan-access { local | radius-scheme radius-scheme-name } *
undo authorization lan-access

```

Views

ISP domain view

Modified command: authorization login

Old syntax

In non-FIPS mode:

```

authorization login { hwtacacs-scheme hwtacacs-scheme-name
[ radius-scheme radius-scheme-name ] [ local ] [ none ] | local [ none ] |
none | radius-scheme radius-scheme-name [ hwtacacs-scheme
hwtacacs-scheme-name ] [ local ] [ none ] }
undo authorization login

```

In FIPS mode:

```

authorization login { hwtacacs-scheme hwtacacs-scheme-name
[ radius-scheme radius-scheme-name ] [ local ] | local | radius-scheme
radius-scheme-name [ hwtacacs-scheme hwtacacs-scheme-name ] [ local ] }
undo authorization login

```

New syntax

In non-FIPS mode:

```

authorization login { hwtacacs-scheme hwtacacs-scheme-name
[ radius-scheme radius-scheme-name ] [ local ] [ none ] | local
[ radius-scheme radius-scheme-name | hwtacacs-scheme
hwtacacs-scheme-name ] * [ none ] | none | radius-scheme
radius-scheme-name [ hwtacacs-scheme hwtacacs-scheme-name ] [ local ]
[ none ] }
undo authorization login

```

In FIPS mode:

```

authorization login { hwtacacs-scheme hwtacacs-scheme-name
[ radius-scheme radius-scheme-name ] [ local ] | local [ radius-scheme
radius-scheme-name | hwtacacs-scheme hwtacacs-scheme-name ] * |
radius-scheme radius-scheme-name [ hwtacacs-scheme hwtacacs-scheme-name ]
[ local ] }
undo authorization login

```

Views

ISP domain view

Modified command: authorization portal

Old syntax

In non-FIPS mode:

```
authorization portal { local [ none ] | none | radius-scheme  
radius-scheme-name [ local ] [ none ] }
```

```
undo authorization portal
```

In FIPS mode:

```
authorization portal { local | radius-scheme radius-scheme-name [ local ] }
```

```
undo authorization portal
```

New syntax

In non-FIPS mode:

```
authorization portal { local [ radius-scheme radius-scheme-name ] [ none ]  
| none | radius-scheme radius-scheme-name [ local ] [ none ] }
```

```
undo authorization portal
```

In FIPS mode:

```
authorization portal { local | radius-scheme radius-scheme-name } *
```

```
undo authorization portal
```

Views

ISP domain view

Modified feature: Setting the 802.1X periodic reauthentication timer

Feature change description

As from this version, the value range for the 802.1X periodic reauthentication timer was changed. The new value range is 60 to 86400 seconds.

Command changes

Modified command: dot1x timer

Syntax

```
dot1x timer { ead-timeout ead-timeout-value | handshake-period  
handshake-period-value | offline-detect offline-detect-value |  
quiet-period quiet-period-value | reauth-period reauth-period-value |  
server-timeout server-timeout-value | supp-timeout supp-timeout-value |  
tx-period tx-period-value | unicast-trigger quiet-period  
quiet-period-value | user-aging { auth-fail-vlan | auth-fail-vsi |  
critical-vlan | critical-vsi | guest-vlan | guest-vsi } aging-time-value }  
  
undo dot1x timer { ead-timeout | handshake-period | offline-detect |  
quiet-period | reauth-period | server-timeout | supp-timeout | tx-period |
```

```
unicast-trigger  quiet-period  |  user-aging  {  auth-fail-vlan  |
auth-fail-vsi | critical-vlan | critical-vsi | guest-vlan | guest-vsi } }
```

Views

System view

Change description

Before modification: The value range for the *reauth-period-value* argument was 60 to 7200 seconds.

After modification: The value range for the *reauth-period-value* argument is 60 to 86400 seconds.

Modified command: dot1x timer reauth-period (interface view)

Syntax

```
dot1x timer reauth-period reauth-period-value
undo dot1x timer reauth-period
```

Views

Layer 2 Ethernet interface view

Layer 2 aggregate interface view

Change description

Before modification: The value range for the *reauth-period-value* argument was 60 to 7200 seconds.

After modification: The value range for the *reauth-period-value* argument is 60 to 86400 seconds.

Modified feature: Setting the periodic MAC reauthentication timer

Feature change description

As from this version, the value range for the periodic MAC reauthentication timer was changed. The new value range is 60 to 86400 seconds.

Command changes

Modified command: mac-authentication timer (interface view)

Syntax

```
mac-authentication timer { auth-delay auth-delay-time | reauth-period
reauth-period-value }
undo mac-authentication timer { auth-delay | reauth-period }
```

Views

Layer 2 Ethernet interface view

Layer 2 aggregate interface view

Change description

Before modification: The value range for the *reauth-period-value* argument was 60 to 7200 seconds.

After modification: The value range for the *reauth-period-value* argument is 60 to 86400 seconds.

Modified command: mac-authentication timer (system view)

Syntax

```
mac-authentication timer { offline-detect offline-detect-value | quiet
quiet-value | reauth-period reauth-period-value | server-timeout
server-timeout-value | temporary-user-aging aging-time-value |
user-aging { critical-vlan | critical-vsi | guest-vlan | guest-vsi }
aging-time-value }

undo mac-authentication timer { offline-detect | quiet | reauth-period |
server-timeout | temporary-user-aging | user-aging { critical-vlan |
critical-vsi | guest-vlan | guest-vsi } }
```

Views

System view

Change description

Before modification: The value range for the *reauth-period-value* argument was 60 to 7200 seconds.

After modification: The value range for the *reauth-period-value* argument is 60 to 86400 seconds.

Modified feature: Enabling SNMP notifications for L3VPN

Feature change description

As from this release, you can configure the SNMP notification types to be generated for MPLS L3VPN.

Command changes

Modified command: snmp-agent trap enable l3vpn

Old syntax

```
snmp-agent trap enable l3vpn
undo snmp-agent trap enable l3vpn
```

New syntax

```
snmp-agent trap enable l3vpn [ vrf-down | vrf-ipv6-up | vrf-ipv6-down |
vrf-up ] *
undo snmp-agent trap enable l3vpn [ vrf-down | vrf-ipv6-up | vrf-ipv6-down
| vrf-up ] *
```

Views

System view

Change description

Keywords were added to the command syntax. You can specify the keywords to enable specific types of SNMP notifications for L3VPN.

The following keywords were added:

- **vrf-down**: VPN instance interface down notification. This notification is generated for a VPN instance in the following conditions:
 - The physical state of all the interfaces bound to the VPN instance is down.
 - All physically up interfaces bound to the VPN instance are unbound from the VPN instance.
- **vrf-ipv6-up**: VPN instance interface IPv6 up notification. This notification is generated for a VPN instance in the following conditions:
 - The IPv6 protocol state of the first interface bound to the VPN instance comes up.
 - When the IPv6 protocol state of all interfaces bound to the VPN instance is Down, one or multiple of the interfaces' IPv6 protocol state changes from Down to Up.
- **vrf-ipv6-down**: VPN instance interface IPv6 down notification. This notification is generated for a VPN instance in the following conditions:
 - The IPv6 protocol state of all the interfaces bound to the VPN instance is down.
 - All IPv6 up interfaces bound to the VPN instance are unbound from the VPN instance.
- **vrf-up**: VPN instance interface up notification. This notification is generated in the following conditions:
 - The physical state of the first interface bound to the VPN instance comes up.
 - When the physical state of all interfaces bound to the VPN instance is Down, one or multiple of the interfaces' physical state changes from Down to Up.

Modified feature: Displaying BGP VPNv4 routing information

Feature change description

As from this release, more types of the BGP VPNv4 routing information can be displayed.

Command changes

Modified command: `display bgp routing-table vpnv4`

Old syntax

```
display bgp [ instance instance-name ] routing-table vpnv4
[ [ route-distinguisher route-distinguisher ] [ ipv4-address
[ { mask-length | mask } [ longest-match ] ] | ipv4-address [ mask-length |
mask ] advertise-info | as-path-acl as-path-acl-number | community-list
{ { basic-community-list-number | comm-list-name } [ whole-match ] |
adv-community-list-number } ] | [ vpn-instance vpn-instance-name ] peer
ipv4-address { advertised-routes | received-routes } [ ipv4-address
[ mask-length | mask ] | statistics ] | statistics ]
```

New syntax

```
display bgp [ instance instance-name ] routing-table vpnv4
[ [ route-distinguisher route-distinguisher ] [ ipv4-address
[ { mask-length | mask } [ longest-match ] ] | ipv4-address [ mask-length |
mask ] advertise-info | ipv4-address [ mask-length | mask ] { as-path |
cluster-list | community | ext-community } | [ vpn-instance
vpn-instance-name ] peer ipv4-address { advertised-routes |
received-routes } [ ipv4-address [ mask-length | mask ] | statistics ] |
statistics ]

display bgp [ instance instance-name ] routing-table vpnv4
[ route-distinguisher route-distinguisher ] as-path-acl
{ as-path-acl-number | as-path-acl-name }

display bgp [ instance instance-name ] routing-table vpnv4
[ route-distinguisher route-distinguisher ] [ statistics ] community
[ community-number<1-32> | aa:nn<1-32> ] [ internet | no-advertise |
no-export | no-export-subconfed ] [ whole-match ]

display bgp [ instance instance-name ] routing-table vpnv4
[ route-distinguisher route-distinguisher ] [ statistics ] community-list
basic-community-list-number | adv-community-list-number |
comm-list-name } [ whole-match ]

display bgp [ instance instance-name ] routing-table vpnv4
[ route-distinguisher route-distinguisher ] [ statistics ] ext-community
[ rt route-target | soo site-of-origin ]<1-32> [ whole-match ]
```

Views

Any view

Change description

Keywords were added to the command syntax. You can specify these keywords to display the corresponding types of BGP VPNv4 route information.

The following keywords were added:

- **as-path**: Displays AS_PATH attribute information for the specified BGP VPNv4 unicast route.
- **cluster-list**: Displays CLUSTER_LIST attribute information for the specified BGP VPNv4 unicast route.
- **community**: Displays COMMUNITY attribute information for the specified BGP VPNv4 unicast route, or displays BGP VPNv4 unicast routes that match the specified community attribute.
- **ext-community**: Displays extended community attribute information for the specified BGP VPNv4 unicast route, or displays BGP VPNv4 unicast routes that match the specified extended community attribute.
- **as-path-acl as-path-acl-name**: Displays BGP VPNv4 routes that match the AS path list specified by its name, a case-sensitive string of 1 to 51 characters. The name cannot contain only digits.
- **community-number<1-32>**: Specifies a community sequence number. The value range for the *community-number* argument is 1 to 4294967295. <1-32> indicates that a maximum of 32 numbers can be specified.
- **aa:nn<1-32>**: Specifies a community number. Both *aa* and *nn* are in the range of 0 to 65535. <1-32> indicates that a maximum of 32 numbers can be specified.
- **internet**: Displays BGP VPNv4 routes that have the INTERNET community attribute. Routes with this attribute can be advertised to all BGP peers. By default, all routes have this attribute.

- **no-advertise**: Displays BGP VPNv4 routes that have the NO_ADVERTISE community attribute. Routes with this attribute cannot be advertised to any peers.
- **no-export**: Displays BGP VPNv4 routes that have the NO_EXPORT community attribute. Routes with this attribute cannot be advertised outside the local AS or confederation, but can be advertised to other sub-ASs in the confederation.
- **no-export-subconfed**: Displays BGP VPNv4 routes that have the NO_EXPORT_SUBCONFED community attribute. Routes with this attribute cannot be advertised outside the local AS or to other sub-ASs in the confederation.
- **rt route-target**: Specifies the Route Target (RT) extended community attribute. An RT is a string of 3 to 24 characters.
- **soo site-of-origin**: Specifies the Site of Origin (SoS) extended community attribute. An SoO is a string of 3 to 24 characters.

An RT or SoS can be in one of the following formats:

- *16-bit AS number.32-bit user-defined number*, for example, 101:3. The value range for the AS number is 0 to 65535. The value range for the user-defined number is 0 to 4294967295.
- *32-bit IP address.16-bit user-defined number*, for example, 192.168.122.15:1. The value range for the user-defined number is 0 to 4294967295.
- *32-bit AS number.16-bit user-defined number*, for example, 70000:3. The value range for the AS number is 65536 to 4294967295. The value range for the user-defined number is 0 to 65535.
- *32-bit IP address/IPv4 mask.16-bit user-defined number*, for example, 192.168.122.15/24:1.
- *32-bit AS number in dotted notation.16-bit user-defined number*, for example, 65535.65535:1.

Modified feature: Displaying BGP VPNv6 routing information

Feature change description

As from this release, more types of the BGP VPNv6 routing information can be displayed.

Command changes

Modified command: display bgp routing-table vpnv6

Old syntax

```
display bgp [ instance instance-name ] routing-table vpnv6
[ [ route-distinguisher route-distinguisher ] [ ipv6-address
prefix-length [ advertise-info ] | as-path-acl as-path-acl-number |
community-list { { basic-community-list-number | comm-list-name }
[ whole-match ] | adv-community-list-number } ] | peer ipv4-address
{ advertised-routes | received-routes } [ ipv6-address prefix-length |
statistics ] | statistics ]
```

New syntax

```
display bgp [ instance instance-name ] routing-table vpnv6
[ [ route-distinguisher route-distinguisher ] [ ipv6-address
prefix-length [ advertise-info ] | ipv6-address prefix-length { as-path |
cluster-list | community | ext-community } | peer ipv4-address
```

```

{ advertised-routes | received-routes } [ ipv6-address prefix-length |
statistics ] | statistics ]

display bgp [ instance instance-name ] routing-table vpnv6
[ route-distinguisher route-distinguisher ] as-path-acl
{ as-path-acl-number | as-path-acl-name }

display bgp [ instance instance-name ] routing-table vpnv6
[ route-distinguisher route-distinguisher ] [ statistics ] community
[ community-number&<1-32> | aa:nn&<1-32> ] [ internet | no-advertise |
no-export | no-export-subconfed ] [ whole-match ]

display bgp [ instance instance-name ] routing-table vpnv6
[ route-distinguisher route-distinguisher ] [ statistics ] community-list
{ basic-community-list-number | adv-community-list-number |
comm-list-name } [ whole-match ]

display bgp [ instance instance-name ] routing-table vpnv6
[ route-distinguisher route-distinguisher ] [ statistics ] ext-community
[ rt route-target | soo site-of-origin ]&<1-32> [ whole-match ]

```

Views

Any view

Change description

Keywords were added to the command syntax. You can specify these keywords to display the corresponding types of BGP VPNv6 route information.

The following keywords were added:

as-path: Displays AS_PATH attribute information for the specified BGP VPNv6 unicast route.

cluster-list: Displays CLUSTER_LIST attribute information for the specified BGP VPNv6 unicast route.

community: Displays COMMUNITY attribute information for the specified BGP VPNv6 unicast route, or displays BGP VPNv6 unicast routes that match the specified community attribute.

ext-community: Displays extended community attribute information for the specified BGP VPNv6 unicast route, or displays BGP VPNv6 unicast routes that match the specified extended community attribute.

as-path-acl as-path-acl-name: Displays BGP VPNv6 routes that match the AS path list specified by its name, a case-sensitive string of 1 to 51 characters. The name cannot contain only digits.

community-number&<1-32>: Specifies a community sequence number. The value range for the *community-number* argument is 1 to 4294967295. &<1-32> indicates that a maximum of 32 numbers can be specified.

aa:nn&<1-32>: Specifies a community number. Both *aa* and *nn* are in the range of 0 to 65535. &<1-32> indicates that a maximum of 32 numbers can be specified.

internet: Displays BGP VPNv6 routes that have the INTERNET community attribute. Routes with this attribute can be advertised to all BGP peers. By default, all routes have this attribute.

no-advertise: Displays BGP VPNv6 routes that have the NO_ADVERTISE community attribute. Routes with this attribute cannot be advertised to any peers.

no-export: Displays BGP VPNv6 routes that have the NO_EXPORT community attribute. Routes with this attribute cannot be advertised outside the local AS or confederation, but can be advertised to other sub-ASs in the confederation.

no-export-subconfed: Displays BGP VPNv6 routes that have the NO_EXPORT_SUBCONFED community attribute. Routes with this attribute cannot be advertised outside the local AS or to other sub-ASs in the confederation.

rt route-target: Specifies the Route Target (RT) extended community attribute. An RT is a string of 3 to 24 characters.

soo site-of-origin: Specifies the Site of Origin (SoS) extended community attribute. An SoO is a string of 3 to 24 characters.

An RT or SoS can be in one of the following formats:

- *16-bit AS number.32-bit user-defined number*, for example, 101:3. The value range for the AS number is 0 to 65535. The value range for the user-defined number is 0 to 4294967295.
- *32-bit IP address.16-bit user-defined number*, for example, 192.168.122.15:1. The value range for the user-defined number is 0 to 4294967295.
- *32-bit AS number.16-bit user-defined number*, for example, 70000:3. The value range for the AS number is 65536 to 4294967295. The value range for the user-defined number is 0 to 65535.
- *32-bit IP address/IPv4 mask.16-bit user-defined number*, for example, 192.168.122.15/24:1.
- *32-bit AS number in dotted notation.16-bit user-defined number*, for example, 65535.65535:1.

Modified feature: Creating an SNMPv3 user

Feature change description

As from this release, you can specify the SHA224, SHA256, SHA384, and SHA512 authentication algorithm for creating an SNMPv3 user. To configure the authentication or encryption key in encrypted form, you can specify the 3DESSHA224, 3DESSHA256, 3DESSHA384, 3DESSHA512, AES192SHA224, AES192SHA256, AES192SHA384, AES192SHA512, AES256SHA224, AES256SHA256, AES256SHA384, AES256SHA512, SHA224, SHA256, SHA384, or SHA512 authentication and encryption algorithm to calculate the key from its plaintext form to encrypted form. The 3DESSHA224, 3DESSHA256, 3DESSHA384, and 3DESSHA512 authentication and encryption algorithms are supported only in non-FIPS mode.

Command changes

Modified command: snmp-agent usm-user v3

Old syntax

In non-FIPS mode:

- In VACM mode:

```
snmp-agent usm-user v3 user-name group-name [ remote { ipv4-address |  
ipv6 ipv6-address } [ vpn-instance vpn-instance-name ] ] [ { cipher |  
simple } authentication-mode { md5 | sha } auth-password [ privacy-mode  
{ 3des | aes128 | aes192 | aes256 | des56 } priv-password ] ] [ acl  
{ ipv4-acl-number | name ipv4-acl-name } | acl ipv6 { ipv6-acl-number  
| name ipv6-acl-name } ] *  
  
undo snmp-agent usm-user v3 user-name { local | engineid  
engineid-string | remote { ipv4-address | ipv6 ipv6-address }  
[ vpn-instance vpn-instance-name ] }
```

- In RBAC mode:

```
snmp-agent usm-user v3 user-name user-role role-name [ remote
{ ipv4-address | ipv6 ipv6-address } [ vpn-instance vpn-instance-name ] ]
[ { cipher | simple } authentication-mode { md5 | sha } auth-password
[ privacy-mode { 3des | aes128 | aes192 | aes256 | des56 } priv-password ] ]
[ acl { ipv4-acl-number | name ipv4-acl-name } | acl ipv6
{ ipv6-acl-number | name ipv6-acl-name } ] *

undo snmp-agent usm-user v3 user-name { local | engineid
engineid-string | remote { ipv4-address | ipv6 ipv6-address }
[ vpn-instance vpn-instance-name ] }
```

In FIPS mode:

- In VACM mode:

```
snmp-agent usm-user v3 user-name group-name [ remote { ipv4-address |
ipv6 ipv6-address } [ vpn-instance vpn-instance-name ] ] { cipher |
simple } authentication-mode sha auth-password [ privacy-mode { aes128
| aes192 | aes256 } priv-password ] [ acl { ipv4-acl-number | name
ipv4-acl-name } | acl ipv6 { ipv6-acl-number | name ipv6-acl-name } ]
*

undo snmp-agent usm-user v3 user-name { local | engineid
engineid-string | remote { ipv4-address | ipv6 ipv6-address }
[ vpn-instance vpn-instance-name ] }
```

- In RBAC mode:

```
snmp-agent usm-user v3 user-name user-role role-name [ remote
{ ipv4-address | ipv6 ipv6-address } [ vpn-instance vpn-instance-name ] ]
[ { cipher | simple } authentication-mode sha auth-password
[ privacy-mode { aes128 | aes192 | aes256 } priv-password ] ] [ acl
{ ipv4-acl-number | name ipv4-acl-name } | acl ipv6 { ipv6-acl-number
| name ipv6-acl-name } ] *

undo snmp-agent usm-user v3 user-name { local | engineid
engineid-string | remote { ipv4-address | ipv6 ipv6-address }
[ vpn-instance vpn-instance-name ] }
```

New syntax

In non-FIPS mode:

- In VACM mode:

```
snmp-agent usm-user v3 user-name group-name [ remote { ipv4-address |
ipv6 ipv6-address } [ vpn-instance vpn-instance-name ] ] [ { cipher |
simple } authentication-mode { md5 | sha | sha224 | sha256 | sha384 |
sha512 } auth-password [ privacy-mode { 3des | aes128 | aes192 | aes256 |
des56 } priv-password ] ] [ acl { ipv4-acl-number | name ipv4-acl-name }
| acl ipv6 { ipv6-acl-number | name ipv6-acl-name } ] *

undo snmp-agent usm-user v3 user-name { local | engineid
engineid-string | remote { ipv4-address | ipv6 ipv6-address }
[ vpn-instance vpn-instance-name ] }
```

- In RBAC mode:

```
snmp-agent usm-user v3 user-name user-role role-name [ remote
{ ipv4-address | ipv6 ipv6-address } [ vpn-instance vpn-instance-name ] ]
[ { cipher | simple } authentication-mode { md5 | sha | sha224 | sha256 |
sha384 | sha512 } auth-password [ privacy-mode { 3des | aes128 | aes192 |
aes256 | des56 } priv-password ] ] [ acl { ipv4-acl-number | name
```

```

    ipv4-acl-name } | acl ipv6 { ipv6-acl-number | name ipv6-acl-name } ]
    *

    undo snmp-agent usm-user v3 user-name { local | engineid
    engineid-string | remote { ipv4-address | ipv6 ipv6-address }
    [ vpn-instance vpn-instance-name ] }

```

In FIPS mode:

- In VACM mode:

```

    snmp-agent usm-user v3 user-name group-name [ remote { ipv4-address |
    ipv6 ipv6-address } [ vpn-instance vpn-instance-name ] ] { cipher |
    simple } authentication-mode { sha | sha224 | sha256 | sha384 | sha512 }
    auth-password [ privacy-mode { aes128 | aes192 | aes256 } priv-password ]
    [ acl { ipv4-acl-number | name ipv4-acl-name } | acl ipv6
    { ipv6-acl-number | name ipv6-acl-name } ] *

    undo snmp-agent usm-user v3 user-name { local | engineid
    engineid-string | remote { ipv4-address | ipv6 ipv6-address }
    [ vpn-instance vpn-instance-name ] }

```

- In RBAC mode:

```

    snmp-agent usm-user v3 user-name user-role role-name [ remote
    { ipv4-address | ipv6 ipv6-address } [ vpn-instance vpn-instance-name ] ]
    [ { cipher | simple } authentication-mode { sha | sha224 | sha256 | sha384
    | sha512 } auth-password [ privacy-mode { aes128 | aes192 | aes256 }
    priv-password ] ] [ acl { ipv4-acl-number | name ipv4-acl-name } | acl
    ipv6 { ipv6-acl-number | name ipv6-acl-name } ] *

    undo snmp-agent usm-user v3 user-name { local | engineid
    engineid-string | remote { ipv4-address | ipv6 ipv6-address }
    [ vpn-instance vpn-instance-name ] }

```

Views

System view

Change description

Before modification: The SHA224, SHA256, SHA384, and SHA512 authentication algorithms are not supported.

After modification: The SHA224, SHA256, SHA384, and SHA512 authentication algorithms are supported.

Modified command: snmp-agent calculate-password

Old syntax

In non-FIPS mode:

```

    snmp-agent calculate-password plain-password mode { 3desmd5 | 3dessha |
    aes192md5 | aes192sha | aes256md5 | aes256sha | md5 | sha } { local-engineid |
    specified-engineid engineid }

```

In FIPS mode:

```

    snmp-agent calculate-password plain-password mode { aes192sha | aes256sha
    | sha } { local-engineid | specified-engineid engineid }

```

New syntax

In non-FIPS mode:

```
snmp-agent calculate-password plain-password mode { 3desmd5 | 3dessha |
3dessha224 | 3dessha256 | 3dessha384 | 3dessha512 | aes192md5 | aes192sha |
aes192sha224 | aes192sha256 | aes192sha384 | aes192sha512 | aes256md5 |
aes256sha | aes256sha224 | aes256sha256 | aes256sha384 | aes256sha512 | md5
| sha | sha224 | sha256 | sha384 | sha512 } { local-engineid |
specified-engineid engineid }
```

In FIPS mode:

```
snmp-agent calculate-password plain-password mode { aes192sha |
aes192sha224 | aes192sha256 | aes192sha384 | aes192sha512 | aes256sha |
aes256sha224 | aes256sha256 | aes256sha384 | aes256sha512 | sha | sha224 |
sha256 | sha384 | sha512 } { local-engineid | specified-engineid engineid }
```

Views

System view

Change description

Before modification:

- *plain-password*: Specifies an plaintext-form key. The argument is a case-sensitvie string of 1 to 64 characters.
- The 3DESSHA224, 3DESSHA256, 3DESSHA384, 3DESSHA512, AES192SHA224, AES192SHA256, AES192SHA384, AES192SHA512, AES256SHA224, AES256SHA256, AES256SHA384, AES256SHA512, SHA224, SHA256, SHA384, and SHA512 authentication and encryption algorithms are not supported in non-FIPS mode.
- The AES192SHA224, AES192SHA256, AES192SHA384, AES192SHA512, AES256SHA224, AES256SHA256, AES256SHA384, AES256SHA512, SHA224, SHA256, SHA384, and SHA512 authentication and encryption algorithms are not supported in FIPS mode.

After modification:

- *plain-password*: Specifies an plaintext-form key. The argument is a case-sensitvie string of 1 to 128 characters.
- The 3DESSHA224, 3DESSHA256, 3DESSHA384, 3DESSHA512, AES192SHA224, AES192SHA256, AES192SHA384, AES192SHA512, AES256SHA224, AES256SHA256, AES256SHA384, AES256SHA512, SHA224, SHA256, SHA384, and SHA512 authentication and encryption algorithms are supported in non-FIPS mode.
- The AES192SHA224, AES192SHA256, AES192SHA384, AES192SHA512, AES256SHA224, AES256SHA256, AES256SHA384, AES256SHA512, SHA224, SHA256, SHA384, and SHA512 authentication and encryption algorithms are supported in FIPS mode.

Modified feature: Displaying local public keys

Feature change description

The **Key length** field was added to the output for the **display public-key local public** command to indicate the key length of public keys in local key pairs, in bits.

Command changes

Modified command: display public-key local public

Syntax

```
display public-key local { dsa | ecdsa | rsa } public [ name key-name ]
```

Views

Any view

Change description

The **Key length** field was added to the output for this command to indicate the key length of public keys in local key pairs, in bits.

Modified feature: Flow-mirroring traffic to an interface

Feature change description

When you configure flow-mirroring traffic to an interface, you can configure the destination interface to send the mirrored packets to the specified reflector port. Then, the reflector port broadcasts the mirrored packets within the specified VLAN.

Command changes

Modified command: mirror-to interface

Old syntax

Syntax 1:

```
mirror-to interface interface-type interface-number [ loopback |  
[ destination-ip destination-ip-address source-ip source-ip-address  
[ dscp dscp-value | vlan vlan-id | vrf-instance vrf-name ] * ] ]
```

Syntax 2:

```
mirror-to interface destination-ip destination-ip-address source-ip  
source-ip-address [ dscp dscp-value | vlan vlan-id | vrf-instance vrf-name ]  
*
```

New syntax

Syntax 1:

```
mirror-to interface interface-type interface-number [ loopback |  
[ destination-ip destination-ip-address source-ip source-ip-address  
[ dscp dscp-value | vlan vlan-id | vrf-instance vrf-name ] * ] ]
```

Syntax 2:

```
mirror-to interface destination-ip destination-ip-address source-ip  
source-ip-address [ dscp dscp-value | vlan vlan-id | vrf-instance vrf-name ]  
*
```

Syntax 3:

```
mirror-to interface interface-type interface-number reflector-port  
interface-type interface-number strip-vlan vlan-id
```

Views

Traffic behavior view

Parameters

reflector-port *interface-type interface-number*: Specifies a reflector port by its type and number.

strip-vlan *vlan-id*: Broadcasts mirrored packets within a VLAN specified by its ID in the range of 1 to 4094.

Change description

Syntax 3 was added.

When you use syntax 3 to configure flow-mirroring traffic to an interface, the device copies packets received on the mirroring sources to the destination interface. Then, the destination interface sends the mirrored packets to the specified reflector port, and the reflector port broadcasts the mirrored packets within the specified VLAN. This syntax is applicable only when the mirrored packets are VLAN-tagged and these packets must be sent out of the device untagged. When you use syntax 3, the specified mirroring destination interface and reflector port must be assigned to a mirroring-type service loopback group. For more information about service loopback groups, see service loopback group configuration in *Layer 2—LAN Switching Configuration Guide*.

Modified feature: Disabling the device from learning the source MAC addresses of Layer 2 protocol packets

Feature change description

Before modification: By default, the device learns the source MAC addresses of Layer 2 protocol packets.

After modification: By default, the device does not learn the source MAC addresses of Layer 2 protocol packets. The source MAC addresses cover the MAC addresses of the following Layer 2 protocol packets:

- BPDUs destined for a MAC address in the range of 0x01-80-c2-00-00-00 to 0x01-80-c2-00-00-0f.
- GARP PDUs destined for a MAC address in the range of 0x01-80-c2-00-00-20 to 0x01-80-c2-00-00-2f.
- PVST BPDUs destined for MAC address 0x01-00-0c-cc-cc-cd.

Command changes

Modified command: `mac-address mac-learning pdu`

Syntax

`mac-address mac-learning pdu`

`undo mac-address mac-learning pdu`

.Views

System view

Change description

Before modification: By default, the device learns the source MAC addresses of Layer 2 protocol packets.

After modification: By default, the device does not learn the source MAC addresses of Layer 2 protocol packets.

Modified feature: Specifying multiple VLAN IDs for VLAN name TLVs

Feature change description

Before modification: When you configure a port to advertise VLAN name TLVs, you can specify only one VLAN ID.

After modification: When you configure a port to advertise VLAN name TLVs, you can specify multiple VLAN IDs. This modification applies to Layer 2 Ethernet interfaces and Layer 2 aggregate interfaces.

When you execute the **display lldp local-information** and **display lldp neighbor-information** commands, the command output displays information about all advertised and received VLAN name TLVs.

Command changes

Modified command: lldp tlv-enable

Old syntax

In Layer 2 Ethernet interface view:

- For nearest bridge agents:

```
lldp tlv-enable { basic-tlv { all | port-description |
system-capability | system-description | system-name |
management-address-tlv [ ipv6 ] [ ip-address | interface loopback
interface-number ] } | dot1-tlv { all | port-vlan-id | link-aggregation
| dcbx | protocol-vlan-id [ vlan-id ] | vlan-name [ vlan-id ] |
management-vid [ mvlan-id ] } | dot3-tlv { all | link-aggregation |
mac-physic | max-frame-size | power } | med-tlv { all | capability |
inventory | network-policy [ vlan-id ] | power-over-ethernet |
location-id { civic-address device-type country-code { ca-type
ca-value } &<1-10> | elin-address tel-number } } }

undo lldp tlv-enable { basic-tlv { all | port-description |
system-capability | system-description | system-name |
management-address-tlv [ ipv6 ] [ ip-address | interface loopback
interface-number ] } | dot1-tlv { all | port-vlan-id | link-aggregation
| dcbx | protocol-vlan-id | vlan-name | management-vid } | dot3-tlv { all
| link-aggregation | mac-physic | max-frame-size | power } | med-tlv
{ all | capability | inventory | network-policy [ vlan-id ] |
power-over-ethernet | location-id } }
```
- For nearest non-TPMR bridge agents:

```
lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ ipv6 ] [ ip-address ] } | dot1-tlv
```

```
{ all | port-vlan-id | link-aggregation } | dot3-tlv { all |
link-aggregation } }

lldp tlv-enable dot1-tlv { protocol-vlan-id [ vlan-id ] | vlan-name
[ vlan-id ] | management-vid [ mvlan-id ] }

undo lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ ipv6 ] [ ip-address ] } |
dot1-tlv { all | port-vlan-id | link-aggregation } | dot3-tlv { all |
link-aggregation } }

undo lldp tlv-enable dot1-tlv { protocol-vlan-id | vlan-name |
management-vid }
```

- For nearest customer bridge agents:

```
lldp agent nearest-customer tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ ipv6 ] [ ip-address ] } |
dot1-tlv { all | port-vlan-id | link-aggregation } | dot3-tlv { all |
link-aggregation } }

lldp tlv-enable dot1-tlv { protocol-vlan-id [ vlan-id ] | vlan-name
[ vlan-id ] | management-vid [ mvlan-id ] }

undo lldp agent nearest-customer tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ ipv6 ] [ ip-address ] } |
dot1-tlv { all | port-vlan-id | link-aggregation } | dot3-tlv { all |
link-aggregation } }

undo lldp tlv-enable dot1-tlv { protocol-vlan-id | vlan-name |
management-vid }
```

In Layer 2 aggregate interface view:

```
lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |
management-address-tlv [ ipv6 ] [ ip-address ] | interface loopback
interface-number ] | port-description | system-capability |
system-description | system-name } | dot1-tlv { all | port-vlan-id } }

lldp agent nearest-customer tlv-enable { basic-tlv { all |
management-address-tlv [ ipv6 ] [ ip-address ] | port-description |
system-capability | system-description | system-name } | dot1-tlv { all |
port-vlan-id } }

lldp tlv-enable dot1-tlv { protocol-vlan-id [ vlan-id ] | vlan-name
[ vlan-id ] | management-vid [ mvlan-id ] }

undo lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |
management-address-tlv [ ipv6 ] [ ip-address ] | port-description |
system-capability | system-description | system-name } | dot1-tlv { all |
port-vlan-id } }

undo lldp agent nearest-customer tlv-enable { basic-tlv { all |
management-address-tlv [ ipv6 ] [ ip-address ] | interface loopback
interface-number ] | port-description | system-capability |
system-description | system-name } | dot1-tlv { all | port-vlan-id } }

undo lldp tlv-enable dot1-tlv { protocol-vlan-id | vlan-name |
management-vid }
```

New syntax

In Layer 2 Ethernet interface view:

- For nearest bridge agents:


```
lldp tlv-enable { basic-tlv { all | port-description |
system-capability | system-description | system-name |
management-address-tlv [ ipv6 ] [ ip-address | interface loopback
interface-number ] } | dot1-tlv { all | port-vlan-id | link-aggregation
| dcbbx | protocol-vlan-id [ vlan-id ] | vlan-name [ vlan-id-list ] |
management-vid [ mvlan-id ] } | dot3-tlv { all | link-aggregation |
mac-physic | max-frame-size | power } | med-tlv { all | capability |
inventory | network-policy [ vlan-id ] | power-over-ethernet |
location-id { civic-address device-type country-code { ca-type
ca-value } &<1-10> | elin-address tel-number } } }

undo lldp tlv-enable { basic-tlv { all | port-description |
system-capability | system-description | system-name |
management-address-tlv [ ipv6 ] [ ip-address | interface loopback
interface-number ] } | dot1-tlv { all | port-vlan-id | link-aggregation
| dcbbx | protocol-vlan-id | vlan-name [ vlan-id-list ] | management-vid }
| dot3-tlv { all | link-aggregation | mac-physic | max-frame-size |
power } | med-tlv { all | capability | inventory | network-policy
[ vlan-id ] | power-over-ethernet | location-id } }
```
- For nearest non-TPMR bridge agents:


```
lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ ipv6 ] [ ip-address ] } | dot1-tlv
{ all | port-vlan-id | link-aggregation } | dot3-tlv { all |
link-aggregation } }

lldp tlv-enable dot1-tlv { protocol-vlan-id [ vlan-id ] | vlan-name
[ vlan-id-list ] | management-vid [ mvlan-id ] }

undo lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ ipv6 ] [ ip-address ] } |
dot1-tlv { all | port-vlan-id | link-aggregation } | dot3-tlv { all |
link-aggregation } }

undo lldp tlv-enable dot1-tlv { protocol-vlan-id | vlan-name
[ vlan-id-list ] | management-vid }
```
- For nearest customer bridge agents:


```
lldp agent nearest-customer tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ ipv6 ] [ ip-address ] } |
dot1-tlv { all | port-vlan-id | link-aggregation } | dot3-tlv { all |
link-aggregation } }

lldp tlv-enable dot1-tlv { protocol-vlan-id [ vlan-id ] | vlan-name
[ vlan-id-list ] | management-vid [ mvlan-id ] }

undo lldp agent nearest-customer tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ ipv6 ] [ ip-address ] } |
dot1-tlv { all | port-vlan-id | link-aggregation } | dot3-tlv { all |
link-aggregation } }

undo lldp tlv-enable dot1-tlv { protocol-vlan-id | vlan-name
[ vlan-id-list ] | management-vid }
```

In Layer 2 aggregate interface view:

```
lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |
management-address-tlv [ ipv6 ] [ ip-address | interface loopback
```

```

interface-number ] | port-description | system-capability |
system-description | system-name } | dot1-tlv { all | port-vlan-id } }

lldp agent nearest-customer tlv-enable { basic-tlv { all |
management-address-tlv [ ipv6 ] [ ip-address ] | port-description |
system-capability | system-description | system-name } | dot1-tlv { all |
port-vlan-id } }

lldp tlv-enable dot1-tlv { protocol-vlan-id [ vlan-id ] | vlan-name
[ vlan-id-list ] | management-vid [ mvlan-id ] }

undo lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |
management-address-tlv [ ipv6 ] [ ip-address ] | port-description |
system-capability | system-description | system-name } | dot1-tlv { all |
port-vlan-id } }

undo lldp agent nearest-customer tlv-enable { basic-tlv { all |
management-address-tlv [ ipv6 ] [ ip-address ] | interface loopback
interface-number ] | port-description | system-capability |
system-description | system-name } | dot1-tlv { all | port-vlan-id } }

undo lldp tlv-enable dot1-tlv { protocol-vlan-id | vlan-name
[ vlan-id-list ] | management-vid }

```

Views

Layer 2 Ethernet interface view

Layer 2 aggregate interface view

Change description

Before modification: If you configure a Layer 2 Ethernet interface or Layer 2 aggregate interface to advertise VLAN name TLVs, you can specify only one VLAN.

After modification: If you configure a Layer 2 Ethernet interface or Layer 2 aggregate interface to advertise VLAN name TLVs, you can specify multiple VLAN IDs.

vlan-name [*vlan-id-list*]: Advertises VLAN name TLVs. The *vlan-id-list* argument specifies a VLAN range in the TLVs to be advertised in the format of { *vlan-id1* [to *vlan-id2*] } &<1-10>. The value range for the *vlan-id* argument is 1 to 4094 and the value for *vlan-id2* must be equal to or greater than the value for *vlan-id1*. &<1-10> indicates that you can specify up to 10 VLAN ID ranges. The default value for the *vlan-id-list* argument is the lowest VLAN ID on the port. If you do not specify a VLAN ID and the port is not assigned to any VLAN, the PVID of the port is advertised or the port cancels advertising all VLAN name TLVs.

Modified feature: Restrictions for an interface to join a Layer 2 aggregation group

Feature change description

Before modification:

- An interface can join a Layer 2 aggregation group only when the interface meets the following requirements:
 - It uses default VLAN configuration.
 - Its port isolation configuration is the same as that on the aggregate interface.
- After leaving a Layer 2 aggregation group, an interface restores its VLAN configuration to the default.

After modification:

- An interface cannot join a Layer 2 aggregation group if it has different attribute configurations from the aggregate interface.
- After leaving a Layer 2 aggregation group, an interface retains the VLAN configuration inheriting from the aggregate interface.

Command changes

None.

Modified feature: Applying a QoS policy to an interface

Feature change description

As from this version, you can apply a QoS policy to the outbound direction of a Layer 2 aggregate interface, Layer 3 aggregate interface or subinterface.

Command changes

Modified command: qos apply policy (interface view)

Syntax

```
qos apply [ ipv6-matching | accounting | mirroring | remarking ] policy
policy-name { inbound | outbound } [ share-mode ]

undo qos apply [ ipv6-matching | accounting | mirroring | remarking ]
policy policy-name { inbound | outbound }
```

Views

Interface view

Change description

Before modification: You cannot apply a QoS policy to the outbound direction of a Layer 2 aggregate interface, Layer 3 aggregate interface, or Layer 3 aggregate subinterface.

After modification: You can apply a generic or IPv6-Matching QoS policy to the outbound direction of a Layer 2 aggregate interface, Layer 3 aggregate interface, or Layer 3 aggregate subinterface.

Feature 6705

This release has the following changes:

- New feature: FEC mode
- New feature: Displaying ND entry statistics
- New feature: User IP address conflict SNMP notifications for ARP
- Modified feature: Restrictions for an interface to join a Layer 2 aggregation group
- Modified feature: Configuring MAC address borrowing
- Modified feature: Configuring the types of advertisable TLVs on a port
- Modified feature: Enabling IGMP proxying on an interface
- Modified feature: DRNI term changes
- Modified feature: Using OSPF commands on interfaces
- Modified feature: Configuring a static route
- Modified feature: Configuring an IPv6 static route

New feature: FEC mode

Configuring FEC

About this task

The forward error correction (FEC) feature corrects packet errors to improve transmission quality. It attaches correction information to a packet at the sending end, and corrects error codes generated during transmission at the receiving end based on the correction information. You can set the FEC mode as needed.

Restrictions and guidelines

Make sure you set the same FEC mode for both interfaces of a link.

Procedure

1. Enter system view.
system-view
2. Enter Ethernet interface view.
interface *interface-type interface-number*
3. Set the FEC mode of the Ethernet interface.
port fec mode { auto | none | rs-fec }
By default, the FEC mode of an Ethernet interface is autonegotiated.

Command reference

port fec mode

Use **port fec mode** to set the forward error correction (FEC) mode of an interface.

Use **undo port fec mode** to restore the default.

Syntax

```
port fec mode { auto | none | rs-fec }  
undo port fec mode
```

Default

The FEC mode of an interface is autonegotiated.

Views

100-GE interface view

Predefined user roles

network-admin

Parameters

auto: Specifies the FEC autonegotiation mode.

none: Performs no FEC.

rs-fec: Specifies the RS-FEC mode.

Usage guidelines

The FEC feature corrects packet errors to improve transmission quality. It attaches correction information to a packet at the sending end, and corrects error codes generated during transmission at the receiving end based on the correction information. You can set the FEC mode as needed.

Make sure you set the same FEC mode for both interfaces of a link.

A 100-GE interface not operating at 100 Gbps does not support FEC mode configuration.

Examples

Set the FEC mode of HundredGigE1/0/54 to autonegotiation.

```
<Sysname> system-view
```

```
[Sysname] interface hundredgige 1/0/54
```

```
[Sysname-HundredGigE1/0/54] port fec mode auto
```

New feature: Displaying ND entry statistics

Displaying ND entry statistics

As from this release, the device supports displaying ND entry statistics.

Command reference

display ipv6 neighbors statistics

Use **display ipv6 neighbors statistics** to display ND entry statistics.

Syntax

```
display ipv6 neighbors statistics { [ by-slot ] all | interface  
{ interface-name / interface-type interface-number } | slot slot-number }
```

Views

Any view

Predefined user roles

network-admin
network-operator

Parameters

all: Displays statistics about all ND entries.

interface interface-type interface-number: Specifies an interface by its type and number.

by-slot: Displays ND entry statistics on a per member device basis.

slot slot-number: Specifies an IRF member device by its member ID. If you do not specify a member device, this command displays ND entry statistics for all member devices.

Usage guidelines

Use ND entry statistics to monitor the usage of entry resources. When an error occurs during packet forwarding, you can view ND entry statistics to identify whether it is because too many entry resources are occupied.

Examples

Display ND entry statistics on Ten-GigabitEthernet 1/0/1.

```
<Sysname> display ipv6 neighbors statistics interface ten-gigabitethernet 1/0/1
```

| State | Dynamic | Static | Rule |
|-------|---------|--------|------|
| Incmp | 0 | 0 | 0 |
| Reach | 0 | 2 | 0 |
| Stale | 1 | - | - |
| Delay | 0 | - | - |
| Probe | 0 | - | - |
| Total | 1 | 2 | 0 |

Display statistics about all ND entries.

```
<Sysname> display ipv6 neighbors statistics all
```

| State | Dynamic | Static | Rule |
|-------|---------|--------|------|
| Incmp | 0 | 4 | 0 |
| Reach | 1 | 2 | 0 |
| Stale | 0 | - | - |
| Delay | 0 | - | - |
| Probe | 0 | - | - |
| Total | 1 | 6 | 0 |

Table 1 Command output

| Field | Description |
|---------|---|
| Dynamic | Number of ND entries obtained dynamically. |
| Static | Number of ND entries configured statically. |
| Rule | Number of ND entries obtained from the IPoE or Portal module. |

| Field | Description |
|-------|---|
| Incmp | Number of ND entries in Incmp state. |
| Reach | Number of ND entries in Reach state. |
| Stale | Number of ND entries in Stale state. |
| Delay | Number of ND entries in Delay state. |
| Probe | Number of ND entries in Probe state. |

New feature: User IP address conflict SNMP notifications for ARP

Enabling user IP address conflict SNMP notifications for ARP

About this task

When a user IP address conflict occurs, this feature enables the device to send a notification to the SNMP module. The notification includes the sender IP address and sender MAC address in the conflicted ARP packet and the MAC address in the corresponding local ARP entry.

For user IP address conflict notifications to be sent correctly, you must also configure SNMP on the device. For more information about SNMP configuration, see *Network Management and Monitoring Configuration Guide*.

Procedure

1. Enter system view.
system-view
2. Enable user IP address conflict SNMP notifications for ARP.
snmp-agent trap enable arp user-ip-conflict
By default, user IP address conflict SNMP notifications for ARP are disabled.

Command reference

Modified command: snmp-agent trap enable arp

Old syntax

```
snmp-agent trap enable arp [ rate-limit ]
undo snmp-agent trap enable arp [ rate-limit ]
```

New syntax

```
snmp-agent trap enable arp [ rate-limit | user-ip-conflict ] *
undo snmp-agent trap enable arp [ rate-limit | user-ip-conflict ] *
```

Views

System view

Change description

The **user-ip-conflict** keyword was added. This keyword enables user IP address conflict SNMP notifications for ARP. If you do not specify a keyword in this command, this command enables all types of SNMP notifications for ARP.

Modified feature: Restrictions for an interface to join a Layer 2 aggregation group

Feature change description

Before modification:

- An interface cannot join a Layer 2 aggregation group if it has different attribute configurations from the aggregate interface.
- After leaving a Layer 2 aggregation group, an interface retains the VLAN configuration inheriting from the aggregate interface.

After modification:

- An interface can join a Layer 2 aggregation group only when the interface meets the following requirements:
 - It uses default VLAN configuration.
 - Its port isolation configuration is the same as that on the aggregate interface.
- After leaving a Layer 2 aggregation group, an interface restores its VLAN configuration to the default.

Command changes

None.

Modified feature: Configuring MAC address borrowing

Feature change description

As from this version, the device can generate an ARP or ND entry after receiving an LLDP frame containing a management address TLV on a Layer 2 aggregate interface. You can also set the source MAC address of LLDP frames in Layer 2 aggregate interface view.

Command changes

Modified command: `lldp management-address`

Old syntax

Layer 2 Ethernet interface view:

```
lldp management-address { arp-learning | nd-learning } vlan vlan-id  
undo lldp management-address { arp-learning | nd-learning }
```

Layer 3 Ethernet interface view:

```
lldp management-address { arp-learning | nd-learning } [ vlan vlan-id ]  
undo lldp management-address { arp-learning | nd-learning }
```

New syntax

Layer 2 Ethernet interface view/Layer 2 aggregate interface view:

```
lldp management-address { arp-learning | nd-learning } vlan vlan-id
```

```
undo lldp management-address { arp-learning | nd-learning }
```

Layer 3 Ethernet interface view:

```
lldp management-address { arp-learning | nd-learning } [ vlan vlan-id ]
```

```
undo lldp management-address { arp-learning | nd-learning }
```

Views

Layer 2 Ethernet interface view

Layer 2 aggregate interface view

Layer 3 Ethernet interface view

Change description

Before modification: You cannot enable generation of ARP or ND entries for received management address TLVs on a Layer 2 aggregate interface.

After modification: You can enable generation of ARP or ND entries for received management address TLVs on a Layer 2 aggregate interface.

Modified command: `lldp source-mac vlan`

Syntax

```
lldp source-mac vlan vlan-id
```

```
undo lldp source-mac vlan
```

Views

Layer 2 Ethernet interface view

Layer 2 aggregate interface view

Layer 3 Ethernet interface view

Change description

Before modification: The device does not support setting the source MAC address of outgoing LLDP frames on a Layer 2 aggregate interface.

After modification: The device supports setting the source MAC address of outgoing LLDP frames on a Layer 2 aggregate interface.

Modified feature: Configuring the types of advertisable TLVs on a port

Feature change description

As from this version, the `interface loopback interface-number` option is supported in Layer 2 Ethernet interface view when you configure the types of advertisable TLVs. That is, you can specify the IP address of a loopback interface as the management address in Layer 2 Ethernet interface view.

Command changes

Modified command: lldp tlv-enable

Old syntax

In Layer 2 Ethernet interface view:

- For nearest bridge agents:

```
lldp tlv-enable { basic-tlv { all | port-description |  
system-capability | system-description | system-name |  
management-address-tlv [ ipv6 ] [ ip-address | interface loopback  
interface-number ] } | dot1-tlv { all | port-vlan-id | link-aggregation  
| dcbbx | protocol-vlan-id [ vlan-id ] | vlan-name [ vlan-id ] |  
management-vid [ mvlan-id ] } | dot3-tlv { all | link-aggregation |  
mac-physic | max-frame-size | power } | med-tlv { all | capability |  
inventory | network-policy [ vlan-id ] | power-over-ethernet |  
location-id { civic-address device-type country-code { ca-type  
ca-value } &<1-10> | elin-address tel-number } } }  
  
undo lldp tlv-enable { basic-tlv { all | port-description |  
system-capability | system-description | system-name |  
management-address-tlv [ ipv6 ] [ ip-address | interface loopback  
interface-number ] } | dot1-tlv { all | port-vlan-id | link-aggregation  
| dcbbx | protocol-vlan-id | vlan-name | management-vid } | dot3-tlv { all  
| link-aggregation | mac-physic | max-frame-size | power } | med-tlv  
{ all | capability | inventory | network-policy [ vlan-id ] |  
power-over-ethernet | location-id } }
```
- For nearest non-TPMR bridge agents:

```
lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |  
port-description | system-capability | system-description |  
system-name | management-address-tlv [ ipv6 ] [ ip-address ] } | dot1-tlv  
{ all | port-vlan-id | link-aggregation } | dot3-tlv { all |  
link-aggregation } }  
  
lldp tlv-enable dot1-tlv { protocol-vlan-id [ vlan-id ] | vlan-name  
[ vlan-id ] | management-vid [ mvlan-id ] }  
  
undo lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |  
port-description | system-capability | system-description |  
system-name | management-address-tlv [ ipv6 ] [ ip-address ] } |  
dot1-tlv { all | port-vlan-id | link-aggregation } | dot3-tlv { all |  
link-aggregation } }  
  
undo lldp tlv-enable dot1-tlv { protocol-vlan-id | vlan-name |  
management-vid }
```
- For nearest customer bridge agents:

```
lldp agent nearest-customer tlv-enable { basic-tlv { all |  
port-description | system-capability | system-description |  
system-name | management-address-tlv [ ipv6 ] [ ip-address ] } |  
dot1-tlv { all | port-vlan-id | link-aggregation } | dot3-tlv { all |  
link-aggregation } }  
  
lldp tlv-enable dot1-tlv { protocol-vlan-id [ vlan-id ] | vlan-name  
[ vlan-id ] | management-vid [ mvlan-id ] }  
  
undo lldp agent nearest-customer tlv-enable { basic-tlv { all |  
port-description | system-capability | system-description |  
system-name | management-address-tlv [ ipv6 ] [ ip-address ] } |
```

```

dot1-tlv { all | port-vlan-id | link-aggregation } | dot3-tlv { all |
link-aggregation } }

undo lldp tlv-enable dot1-tlv { protocol-vlan-id | vlan-name |
management-vid }

```

In Layer 3 Ethernet interface view:

```

lldp tlv-enable { basic-tlv { all | port-description | system-capability |
system-description | system-name | management-address-tlv [ ipv6 ]
[ ip-address | interface loopback interface-number ] } | dot1-tlv { all |
link-aggregation } | dot3-tlv { all | link-aggregation | mac-physic |
max-frame-size | power } | med-tlv { all | capability | inventory |
power-over-ethernet | location-id { civic-address device-type
country-code { ca-type ca-value } &<1-10> | elin-address tel-number } } }

lldp agent { nearest-nontpmr | nearest-customer } tlv-enable { basic-tlv
{ all | port-description | system-capability | system-description |
system-name | management-address-tlv [ ipv6 ] [ ip-address ] } | dot1-tlv
{ all | link-aggregation } | dot3-tlv { all | link-aggregation } }

undo lldp tlv-enable { basic-tlv { all | port-description |
system-capability | system-description | system-name |
management-address-tlv [ ipv6 ] [ ip-address | interface loopback
interface-number ] } | dot1-tlv { all | link-aggregation } | dot3-tlv { all |
link-aggregation | mac-physic | max-frame-size | power } | med-tlv { all |
capability | inventory | power-over-ethernet | location-id } }

undo lldp agent { nearest-nontpmr | nearest-customer } tlv-enable
{ basic-tlv { all | port-description | system-capability |
system-description | system-name | management-address-tlv [ ipv6 ]
[ ip-address ] } | dot1-tlv { all | link-aggregation } | dot3-tlv { all |
link-aggregation } }

```

In management Ethernet interface view:

```

lldp tlv-enable { basic-tlv { all | port-description | system-capability |
system-description | system-name | management-address-tlv [ ipv6 ]
[ ip-address ] } | dot1-tlv { all | link-aggregation } | dot3-tlv { all |
link-aggregation | mac-physic | max-frame-size | power } | med-tlv { all |
capability | inventory | power-over-ethernet | location-id { civic-address
device-type country-code { ca-type ca-value } &<1-10> | elin-address
tel-number } } }

lldp agent { nearest-nontpmr | nearest-customer } tlv-enable { basic-tlv
{ all | port-description | system-capability | system-description |
system-name | management-address-tlv [ ipv6 ] [ ip-address ] } | dot1-tlv
{ all | link-aggregation } | dot3-tlv { all | link-aggregation } }

undo lldp tlv-enable { basic-tlv { all | port-description |
system-capability | system-description | system-name |
management-address-tlv [ ipv6 ] [ ip-address ] } | dot1-tlv { all |
link-aggregation } | dot3-tlv { all | link-aggregation | mac-physic |
max-frame-size | power } | med-tlv { all | capability | inventory |
power-over-ethernet | location-id } }

undo lldp agent { nearest-nontpmr | nearest-customer } tlv-enable
{ basic-tlv { all | port-description | system-capability |
system-description | system-name | management-address-tlv [ ipv6 ]
[ ip-address ] } | dot1-tlv { all | link-aggregation } | dot3-tlv { all |
link-aggregation } }

```

In Layer 2 aggregate interface view:

```

lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |
management-address-tlv [ ipv6 ] [ ip-address ] | port-description |
system-capability | system-description | system-name } | dot1-tlv { all |
port-vlan-id } }

lldp agent nearest-customer tlv-enable { basic-tlv { all |
management-address-tlv [ ipv6 ] [ ip-address ] | port-description |
system-capability | system-description | system-name } | dot1-tlv { all |
port-vlan-id } }

lldp tlv-enable dot1-tlv { protocol-vlan-id [ vlan-id ] | vlan-name
[ vlan-id ] | management-vid [ mvlan-id ] }

undo lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |
management-address-tlv [ ipv6 ] [ ip-address ] | port-description |
system-capability | system-description | system-name } | dot1-tlv { all |
port-vlan-id } }

undo lldp agent nearest-customer tlv-enable { basic-tlv { all |
management-address-tlv [ ipv6 ] [ ip-address ] | port-description |
system-capability | system-description | system-name } | dot1-tlv { all |
port-vlan-id } }

undo lldp tlv-enable dot1-tlv { protocol-vlan-id | vlan-name |
management-vid }

```

In Layer 3 aggregate interface view:

```

lldp agent { nearest-customer | nearest-nontpmr } tlv-enable basic-tlv
{ all | management-address-tlv [ ipv6 ] [ ip-address ] | port-description |
system-capability | system-description | system-name }

undo lldp agent { nearest-customer | nearest-nontpmr } tlv-enable basic-tlv
{ all | management-address-tlv [ ipv6 ] [ ip-address ] | port-description |
system-capability | system-description | system-name }

```

In IRF physical interface view:

```

lldp tlv-enable basic-tlv { port-description | system-capability |
system-description | system-name }

undo lldp tlv-enable basic-tlv { port-description | system-capability |
system-description | system-name }

```

New syntax

In Layer 2 Ethernet interface view:

- For nearest bridge agents:

```

lldp tlv-enable { basic-tlv { all | port-description |
system-capability | system-description | system-name |
management-address-tlv [ ipv6 ] [ ip-address | interface loopback
interface-number ] } | dot1-tlv { all | congestion-notification |
port-vlan-id | link-aggregation | dcbx | protocol-vlan-id [ vlan-id ] |
vlan-name [ vlan-id ] | management-vid [ mvlan-id ] } | dot3-tlv { all |
link-aggregation | mac-physic | max-frame-size | power | eee } | med-tlv
{ all | capability | inventory | network-policy [ vlan-id ] |
power-over-ethernet | location-id { civic-address device-type
country-code { ca-type ca-value } &<1-10> | elin-address
tel-number } } }

undo lldp tlv-enable { basic-tlv { all | port-description |
system-capability | system-description | system-name |
management-address-tlv [ ipv6 ] [ ip-address | interface loopback

```

- ```

interface-number] } | dot1-tlv { all | congestion-notification |
port-vlan-id | link-aggregation | dcbx | protocol-vlan-id | vlan-name |
management-vid } | dot3-tlv { all | link-aggregation | mac-physic |
max-frame-size | power | eee } | med-tlv { all | capability | inventory |
network-policy [vlan-id] | power-over-ethernet | location-id } }

```
- For nearest non-TPMR bridge agents:

```

lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ipv6] [ip-address] } | dot1-tlv
{ all | congestion-notification | evb | port-vlan-id | link-aggregation }
| dot3-tlv { all | link-aggregation } }

lldp tlv-enable dot1-tlv { protocol-vlan-id [vlan-id] | vlan-name
[vlan-id] | management-vid [mvlan-id] }

undo lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ipv6] [ip-address] } |
dot1-tlv { all | congestion-notification | evb | port-vlan-id |
link-aggregation } | dot3-tlv { all | link-aggregation } }

undo lldp tlv-enable dot1-tlv { protocol-vlan-id | vlan-name |
management-vid }

```
  - For nearest customer bridge agents:

```

lldp agent nearest-customer tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ipv6] [ip-address] } |
dot1-tlv { all | congestion-notification | port-vlan-id |
link-aggregation } | dot3-tlv { all | link-aggregation } }

lldp tlv-enable dot1-tlv { protocol-vlan-id [vlan-id] | vlan-name
[vlan-id] | management-vid [mvlan-id] }

undo lldp agent nearest-customer tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ipv6] [ip-address] } |
dot1-tlv { all | congestion-notification | port-vlan-id |
link-aggregation } | dot3-tlv { all | link-aggregation } }

undo lldp tlv-enable dot1-tlv { protocol-vlan-id | vlan-name |
management-vid }

```

In Layer 3 Ethernet interface view:

```

lldp tlv-enable { basic-tlv { all | port-description | system-capability |
system-description | system-name | management-address-tlv [ipv6]
[ip-address | interface loopback interface-number] } | dot1-tlv { all |
link-aggregation } | dot3-tlv { all | link-aggregation | mac-physic |
max-frame-size | power | eee } | med-tlv { all | capability | inventory |
power-over-ethernet | location-id { civic-address device-type
country-code { ca-type ca-value } &<1-10> | elin-address tel-number } } }

lldp agent { nearest-nontpmr | nearest-customer } tlv-enable { basic-tlv
{ all | port-description | system-capability | system-description |
system-name | management-address-tlv [ipv6] [ip-address] } | dot1-tlv
{ all | link-aggregation } | dot3-tlv { all | link-aggregation } }

undo lldp tlv-enable { basic-tlv { all | port-description |
system-capability | system-description | system-name |
management-address-tlv [ipv6] [ip-address | interface loopback
interface-number] } | dot1-tlv { all | link-aggregation } | dot3-tlv { all |

```

```

link-aggregation | mac-physic | max-frame-size | power | eee } | med-tlv
{ all | capability | inventory | power-over-ethernet | location-id } }

undo lldp agent { nearest-nontpmr | nearest-customer } tlv-enable
{ basic-tlv { all | port-description | system-capability |
system-description | system-name | management-address-tlv [ipv6]
[ip-address] } | dot1-tlv { all | link-aggregation } | dot3-tlv { all |
link-aggregation } }

```

In management Ethernet interface view:

```

lldp tlv-enable { basic-tlv { all | port-description | system-capability |
system-description | system-name | management-address-tlv [ipv6]
[ip-address] } | dot1-tlv { all | link-aggregation } | dot3-tlv { all |
link-aggregation | mac-physic | max-frame-size | power | eee } | med-tlv
{ all | capability | inventory | power-over-ethernet | location-id
{ civic-address device-type country-code { ca-type ca-value } &<1-10> |
elin-address tel-number } } }

```

```

lldp agent { nearest-nontpmr | nearest-customer } tlv-enable { basic-tlv
{ all | port-description | system-capability | system-description |
system-name | management-address-tlv [ipv6] [ip-address] } | dot1-tlv
{ all | link-aggregation } | dot3-tlv { all | link-aggregation } }

```

```

undo lldp tlv-enable { basic-tlv { all | port-description |
system-capability | system-description | system-name |
management-address-tlv [ipv6] [ip-address] } | dot1-tlv { all |
link-aggregation } | dot3-tlv { all | link-aggregation | mac-physic |
max-frame-size | power | eee } | med-tlv { all | capability | inventory |
power-over-ethernet | location-id } }

```

```

undo lldp agent { nearest-nontpmr | nearest-customer } tlv-enable
{ basic-tlv { all | port-description | system-capability |
system-description | system-name | management-address-tlv [ipv6]
[ip-address] } | dot1-tlv { all | link-aggregation } | dot3-tlv { all |
link-aggregation } }

```

In Layer 2 aggregate interface view:

```

lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |
management-address-tlv [ipv6] [ip-address | interface loopback
interface-number] | port-description | system-capability |
system-description | system-name } | dot1-tlv { all | evb | port-vlan-id } }

```

```

lldp agent nearest-customer tlv-enable { basic-tlv { all |
management-address-tlv [ipv6] [ip-address] | port-description |
system-capability | system-description | system-name } | dot1-tlv { all |
port-vlan-id } }

```

```

lldp tlv-enable dot1-tlv { protocol-vlan-id [vlan-id] | vlan-name
[vlan-id] | management-vid [mvlan-id] }

```

```

undo lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |
management-address-tlv [ipv6] [ip-address] | port-description |
system-capability | system-description | system-name } | dot1-tlv { all |
evb | port-vlan-id } }

```

```

undo lldp agent nearest-customer tlv-enable { basic-tlv { all |
management-address-tlv [ipv6] [ip-address | interface loopback
interface-number] | port-description | system-capability |
system-description | system-name } | dot1-tlv { all | port-vlan-id } }

```

```
undo lldp tlv-enable dot1-tlv { protocol-vlan-id | vlan-name |
management-vid }
```

In Layer 3 aggregate interface view:

```
lldp agent { nearest-customer | nearest-nontpmr } tlv-enable basic-tlv
{ all | management-address-tlv [ipv6] [ip-address] | port-description |
system-capability | system-description | system-name }
```

```
undo lldp agent { nearest-customer | nearest-nontpmr } tlv-enable basic-tlv
{ all | management-address-tlv [ipv6] [ip-address] | port-description |
system-capability | system-description | system-name }
```

In IRF physical interface view:

```
lldp tlv-enable basic-tlv { port-description | system-capability |
system-description | system-name }
```

```
undo lldp tlv-enable basic-tlv { port-description | system-capability |
system-description | system-name }
```

## Views

Layer 2 Ethernet interface view

Layer 3 Ethernet interface view

Management Ethernet interface view

Layer 2 aggregate interface view

Layer 3 aggregate interface view

RF physical interface view

## Change description

Before modification: The **interface loopback** *interface-number* option is not supported in Layer 2 Ethernet interface view.

After modification: The **interface loopback** *interface-number* option is supported in Layer 2 Ethernet interface view.

# Modified feature: Enabling IGMP proxying on an interface

## Feature change description

As from this version, you can enable IGMP proxying on a VSI interface and display multicast routing entries and multicast groups learned through EVPN and maintained by the IGMP proxy.

## Command changes

### Modified command: display igmp group

#### Old syntax

```
display igmp [vpn-instance vpn-instance-name] group [group-address |
interface interface-type interface-number] [static | verbose]
```

## New syntax

```
display igmp [vpn-instance vpn-instance-name] group [evpn]
[group-address | interface interface-type interface-number] [static |
verbose]
```

## Change description

Before modification: The **evpn** keyword is not supported.

After modification: The **evpn** keyword is supported.

**evpn**: Displays multicast groups learned through EVPN. If you do not specify this keyword, the command displays all learned multicast groups.

## Modified command: display igmp proxy group

### Old syntax

```
display igmp [vpn-instance vpn-instance-name] proxy group [group-address
| interface interface-type interface-number] [verbose]
```

### New syntax

```
display igmp [vpn-instance vpn-instance-name] proxy group [evpn]
[group-address | interface interface-type interface-number] [verbose]
```

## Change description

Before modification: The **evpn** keyword is not supported.

After modification: The **evpn** keyword is supported.

**evpn**: Displays multicast routing entries learned through EVPN. If you do not specify this keyword, the command displays all learned multicast routing entries.

## Modified command: igmp proxy enable

### Syntax

```
igmp proxy enable
```

### Views

Interface view

## Change description

Before modification: The command cannot be executed in VSI interface view.

After modification: The command can be executed in VSI interface view.

# Modified feature: DRNI term changes

## Feature change description

The name and terms were changed for the DRNI feature as follows:

- The feature name was changed to M-LAG.
- The **drni** keyword was changed to **m-lag** or **mlag**.
- The **intra-portal-port** keyword was changed to **peer-link**.
- The **ipp** keyword was changed to **peer-link**.

- The **dr** keyword was changed to **m-lag-interface**.
- The terms used in command output and log messages were changed accordingly.

**Table 3 DRNI term changes**

| Old term                | New term                              |
|-------------------------|---------------------------------------|
| DRNI                    | Multichassis link aggregation (M-LAG) |
| DR system               | M-LAG system                          |
| DR interface            | M-LAG interface                       |
| DRNI MAD                | M-LAG MAD                             |
| DR group                | M-LAG group                           |
| DR member device        | M-LAG member device                   |
| DRNI virtual IP address | M-LAG virtual IP address (M-LAG VIP)  |
| IPL                     | M-LAG peer link                       |
| IPP                     | M-LAG peer-link interface             |
| EVPN distributed relay  | EVPN M-LAG                            |

## Command changes in DRNI

### Modified command: display drni consistency

#### Old syntax

```
display drni consistency { type1 | type2 } { global | interface
interface-type interface-number }
```

#### New syntax

```
display m-lag consistency { type1 | type2 } { global | interface
interface-type interface-number }
```

#### Views

Any view

### Modified command: display drni consistency-check status

#### Old syntax

```
display drni consistency-check status
```

#### New syntax

```
display m-lag consistency-check status
```

#### Views

Any view

Modified command: display drni drcp statistics

**Old syntax**

```
display drni drcp statistics [interface interface-type interface-number]
```

**New syntax**

```
display m-lag drcp statistics [interface interface-type interface-number]
```

**Views**

Any view

Modified command: display drni keepalive

**Old syntax**

```
display drni keepalive
```

**New syntax**

```
display m-lag keepalive
```

**Views**

Any view

Modified command: display drni mad verbose

**Old syntax**

```
display drni mad verbose
```

**New syntax**

```
display m-lag mad verbose
```

**Views**

Any view

Modified command: display drni role

**Old syntax**

```
display drni role
```

**New syntax**

```
display m-lag role
```

**Views**

Any view

Modified command: display drni summary

**Old syntax**

```
display drni summary
```

**New syntax**

```
display m-lag summary
```

## Views

Any view

Modified command: display drni system

## Old syntax

```
display drni system
```

## New syntax

```
display m-lag system
```

## Views

Any view

Modified command: display drni troubleshooting

## Old syntax

```
display drni troubleshooting [dr | ipp | keepalive] [history] [count]
```

## New syntax

```
display m-lag troubleshooting [m-lag-interface | peer-link | keepalive]
[history] [count]
```

## Views

Any view

Modified command: display drni verbose

## Old syntax

```
display drni verbose [interface interface-type interface-number]
```

## New syntax

```
display m-lag verbose [interface interface-type interface-number]
```

## Views

Any view

Modified command: display drni virtual-ip

## Old syntax

```
display drni virtual-ip [interface interface-type interface-number]
```

## New syntax

```
display m-lag virtual-ip [interface interface-type interface-number]
```

## Views

Any view

Modified command: drni authentication key

## Old syntax

```
drni authentication key { simple | cipher } string
```

```
undo drni authentication key
```

#### New syntax

```
m-lag authentication key { simple | cipher } string
undo m-lag authentication key
```

#### Views

System view

Modified command: drni auto-recovery reload-delay

#### Old syntax

```
drni auto-recovery reload-delay delay-value
undo drni auto-recovery reload-delay
```

#### New syntax

```
m-lag auto-recovery reload-delay delay-value
undo m-lag auto-recovery reload-delay
```

#### Views

System view

Modified command: drni consistency-check disable

#### Old syntax

```
drni consistency-check disable
undo drni consistency-check disable
```

#### New syntax

```
m-lag consistency-check disable
undo m-lag consistency-check disable
```

#### Views

System view

Modified command: drni consistency-check mode

#### Old syntax

```
drni consistency-check mode { loose | strict }
undo drni consistency-check mode
```

#### New syntax

```
m-lag consistency-check mode { loose | strict }
undo m-lag consistency-check mode
```

#### Views

System view

Modified command: drni drcp period short

**Old syntax**

```
drni drcp period short
undo drni drcp period
```

**New syntax**

```
m-lag drcp period short
undo m-lag drcp period
```

**Views**

Layer 2 aggregate interface view  
Tunnel interface view

Modified command: drni ipp mac-address hold

**Old syntax**

```
drni ipp mac-address hold
undo drni ipp mac-address hold
```

**New syntax**

```
m-lag peer-link mac-address hold
undo m-lag peer-link mac-address hold
```

**Views**

System view

Modified command: drni keepalive { ip | ipv6 }

**Old syntax**

```
drni keepalive { ip | ipv6 } destination { ipv4-address | ipv6-address }
[source { ipv4-address | ipv6-address } | udp-port udp-number |
vpn-instance vpn-instance-name] *
undo drni keepalive { ip | ipv6 }
```

**New syntax**

```
m-lag keepalive { ip | ipv6 } destination { ipv4-address | ipv6-address }
[source { ipv4-address | ipv6-address } | udp-port udp-number |
vpn-instance vpn-instance-name] *
undo m-lag keepalive { ip | ipv6 }
```

**Views**

System view

Modified command: drni keepalive hold-time

**Old syntax**

```
drni keepalive hold-time value
undo drni keepalive hold-time
```

### New syntax

```
m-lag keepalive hold-time value
undo m-lag keepalive hold-time
```

### Views

System view

## Modified command: drni keepalive interval

### Old syntax

```
drni keepalive interval interval [timeout timeout]
undo drni keepalive interval
```

### New syntax

```
m-lag keepalive interval interval [timeout timeout]
undo m-lag keepalive interval
```

### Views

System view

## Modified command: drni mad default-action

### Old syntax

```
drni mad default-action { down | none }
undo drni mad default-action
```

### New syntax

```
m-lag mad default-action { down | none }
undo m-lag mad default-action
```

### Views

System view

## Modified command: drni mad exclude interface

### Old syntax

```
drni mad exclude interface interface-type interface-number
undo drni mad exclude interface interface-type interface-number
```

### New syntax

```
m-lag mad exclude interface interface-type interface-number
undo m-lag mad exclude interface interface-type interface-number
```

### Views

System view

Modified command: drni mad exclude logical-interfaces

**Old syntax**

```
drni mad exclude logical-interfaces
undo drni mad exclude logical-interfaces
```

**New syntax**

```
m-lag mad exclude logical-interfaces
undo m-lag mad exclude logical-interfaces
```

**Views**

System view

Modified command: drni mad include interface

**Old syntax**

```
drni mad include interface interface-type interface-number
undo drni mad include interface interface-type interface-number
```

**New syntax**

```
m-lag mad include interface interface-type interface-number
undo m-lag mad include interface interface-type interface-number
```

**Views**

System view

Modified command: drni mad persistent

**Old syntax**

```
drni mad persistent
undo drni mad persistent
```

**New syntax**

```
m-lag mad persistent
undo m-lag mad persistent
```

**Views**

System view

Modified command: drni mad restore

**Old syntax**

```
drni mad restore
```

**New syntax**

```
m-lag mad restore
```

**Views**

System view

## Modified command: drni restore-delay

### Old syntax

```
drni restore-delay value
undo drni restore-delay
```

### New syntax

```
m-lag restore-delay value
undo m-lag restore-delay
```

### Views

System view

## Modified command: drni role priority

### Old syntax

```
drni role priority priority-value
undo drni role priority
```

### New syntax

```
m-lag role priority priority-value
undo m-lag role priority
```

### Views

System view

## Modified command: drni sequence enable

### Old syntax

```
drni sequence enable
undo drni sequence enable
```

### New syntax

```
m-lag sequence enable
undo m-lag sequence enable
```

### Views

System view

## Modified command: drni standalone enable

### Old syntax

```
drni standalone enable [delay delay-time]
undo drni standalone enable [delay]
```

### New syntax

```
m-lag standalone enable [delay delay-time]
undo m-lag standalone enable [delay]
```

## Views

System view

Modified command: drni system-mac

### Old syntax

```
drni system-mac mac-address
undo drni system-mac
```

### New syntax

```
m-lag system-mac mac-address
undo m-lag system-mac
```

## Views

System view

Modified command: drni system-number

### Old syntax

```
drni system-number system-number
undo drni system-number
```

### New syntax

```
m-lag system-number system-number
undo m-lag system-number
```

## Views

System view

Modified command: drni system-priority

### Old syntax

```
drni system-priority priority
undo drni system-priority
```

### New syntax

```
m-lag system-priority priority
undo m-lag system-priority
```

## Views

System view

Modified command: port drni group

### Old syntax

```
port drni group group-id [allow-single-member]
undo port drni group
```

## New syntax

```
port m-lag group group-id [allow-single-member]
undo port m-lag group
```

## Views

Layer 2 aggregate interface view

## Modified command: port drni intra-portal-port

## Old syntax

```
port drni intra-portal-port port-id
undo port drni intra-portal-port
```

## New syntax

```
port m-lag peer-link port-id
undo port m-lag peer-link
```

## Views

Layer 2 aggregate interface view

Tunnel interface view

## Modified command: port drni ipv6 virtual-ip

## Old syntax

VLAN interface view:

```
port drni ipv6 virtual-ip ipv6-address { prefix-length [active | standby]
[virtual-mac mac-address] | link-local }
undo port drni ipv6 virtual-ip [ipv6-address]
```

Loopback interface view:

```
port drni ipv6 virtual-ip ipv6-address { prefix-length [active | standby]
| link-local }
undo port drni ipv6 virtual-ip [ipv6-address]
```

VSI interface view:

```
port drni ipv6 virtual-ip ipv6-address { prefix-length [active | standby]
| link-local }
undo port drni ipv6 virtual-ip [ipv6-address]
```

## New syntax

VLAN interface view:

```
port m-lag ipv6 virtual-ip ipv6-address { prefix-length [active | standby]
[virtual-mac mac-address] | link-local }
undo port m-lag ipv6 virtual-ip [ipv6-address]
```

Loopback interface view:

```
port m-lag ipv6 virtual-ip ipv6-address { prefix-length [active | standby]
| link-local }
undo port m-lag ipv6 virtual-ip [ipv6-address]
```

VSI interface view:

```
port m-lag ipv6 virtual-ip ipv6-address { prefix-length [active | standby]
| link-local }
undo port m-lag ipv6 virtual-ip [ipv6-address]
```

## Views

VLAN interface view

Loopback interface view

VSI interface view

## Modified command: port drni system-mac

### Old syntax

```
port drni system-mac mac-address
undo port drni system-mac
```

### New syntax

```
port m-lag system-mac mac-address
undo port m-lag system-mac
```

## Views

Aggregate interface view

## Modified command: port drni system-priority

### Old syntax

```
port drni system-priority priority
undo port drni system-priority
```

### New syntax

```
port m-lag system-priority priority
undo port m-lag system-priority
```

## Views

Aggregate interface view

## Modified command: port drni virtual-ip

### Old syntax

VLAN interface view:

```
port drni virtual-ip ipv4-address { mask-length | mask } [active | standby]
virtual-mac mac-address
undo port drni virtual-ip [ipv4-address]
```

Loopback interface view:

```
port drni virtual-ip ipv4-address { mask-length | mask } [active |
standby]
undo port drni virtual-ip [ipv4-address]
```

VSI interface view:

```
port drni virtual-ip ipv4-address { mask-length | mask } [active | standby]
```

```
undo port drni virtual-ip [ipv4-address]
```

## New syntax

VLAN interface view:

```
port m-lag virtual-ip ipv4-address { mask-length | mask } [active | standby] virtual-mac mac-address
```

```
undo port m-lag virtual-ip [ipv4-address]
```

Loopback interface view:

```
port m-lag virtual-ip ipv4-address { mask-length | mask } [active | standby]
```

```
undo port m-lag virtual-ip [ipv4-address]
```

VSI interface view:

```
port m-lag virtual-ip ipv4-address { mask-length | mask } [active | standby]
```

```
undo port m-lag virtual-ip [ipv4-address]
```

## Views

VLAN interface view

Loopback interface view

VSI interface view

## Modified command: reset drni drcp statistics

### Old syntax

```
reset drni drcp statistics [interface interface-list]
```

### New syntax

```
reset m-lag drcp statistics [interface interface-list]
```

## Views

User view

## Modified command: reset drni troubleshooting history

### Old syntax

```
reset drni troubleshooting history
```

### New syntax

```
reset m-lag troubleshooting history
```

## Views

User view

## Command changes in Track

Modified command: track drni-mad-status

### Old syntax

```
track track-entry-number drni-mad-status
undo track track-entry-number
```

### New syntax

```
track track-entry-number mlag-mad-status
undo track track-entry-number
```

### Views

System view

## Command changes in portal

Modified command: portal drni load-sharing-mode

### Old syntax

```
portal drni load-sharing-mode { centralized | distributed { even-ip | odd-ip } }
undo portal drni load-sharing-mode
```

### New syntax

```
portal m-lag load-sharing-mode { centralized | distributed { even-ip | odd-ip } }
undo portal m-lag load-sharing-mode
```

### Views

System view

### Change description

The **drni** keyword was changed to the **m-lag** keyword.

Modified command: portal drni traffic backup

### Old syntax

```
portal drni traffic backup { interval interval-value | threshold threshold-value } *
undo portal drni traffic backup
```

### New syntax

```
portal m-lag traffic backup { interval interval-value | threshold threshold-value } *
undo portal m-lag traffic backup
```

### Views

System view

### Change description

The **drni** keyword was changed to the **m-lag** keyword.

## Command changes in Web authentication

Modified command: display web-auth user

### Old syntax

```
display web-auth user [drni [local | peer]] [interface interface-type
interface-number | slot slot-number]
```

### New syntax

```
display web-auth user [m-lag [local | peer]] [interface interface-type
interface-number | slot slot-number]
```

### Views

Any view

### Change description

The **drni** keyword was changed to the **m-lag** keyword.

## Command changes in AAA

Modified command: nas-ip (RADIUS scheme view)

### Old syntax

```
nas-ip [drni { local | peer }] { ipv4-address | interface interface-type
interface-number | ipv6 ipv6-address }
undo nas-ip [drni { local | peer }] [interface | ipv6]
```

### New syntax

```
nas-ip [m-lag { local | peer }] { ipv4-address | interface interface-type
interface-number | ipv6 ipv6-address }
undo nas-ip [m-lag { local | peer }] [interface | ipv6]
```

### Views

RADIUS scheme view

### Change description

The **drni** keyword was changed to the **m-lag** keyword.

## Command changes in 802.1X

Modified command: display dot1x connection

### Old syntax

```
display dot1x connection [open] [[drni [local | peer]] [interface
interface-type interface-number | online-type { auth-fail-domain |
critical-domain | preauth-domain | success } | slot slot-number | user-name
name-string] | user-mac mac-address]
```

### New syntax

```
display dot1x connection [open] [[m-lag [local | peer]] [interface
interface-type interface-number | online-type { auth-fail-domain |
```

```
critical-domain | preauth-domain | success } | slot slot-number | user-name
name-string] | user-mac mac-address]
```

## Views

Any view

## Change description

The **drni** keyword was changed to the **m-lag** keyword.

# Command changes in MAC authentication

## Modified command: display mac-authentication connection

### Old syntax

```
display mac-authentication connection [open] [[drni [local | peer]]
[interface interface-type interface-number | online-type
{ auth-fail-domain | critical-domain | preauth-domain | success |
url-unavailable-domain } | slot slot-number | user-name user-name] |
user-mac mac-address]
```

### New syntax

```
display mac-authentication connection [open] [[m-lag [local | peer]]
[interface interface-type interface-number | online-type
{ auth-fail-domain | critical-domain | preauth-domain | success |
url-unavailable-domain } | slot slot-number | user-name user-name] |
user-mac mac-address]
```

## Views

Any view

## Change description

The **drni** keyword was changed to the **m-lag** keyword.

# Command changes in port security

## Modified command: display port-security access-user

### Old syntax

```
display port-security access-user [drni [local | peer]] [access-type
{ dot1x | mac-auth | web-auth | static } | domain domain-name | online-type
{ auth-fail-domain | critical-domain | preauth-domain | success |
url-unavailable-domain } | slot slot-number] *
```

### New syntax

```
display port-security access-user [m-lag [local | peer]] [access-type
{ dot1x | mac-auth | web-auth | static } | domain domain-name | online-type
{ auth-fail-domain | critical-domain | preauth-domain | success |
url-unavailable-domain } | slot slot-number] *
```

## Views

Any view

## Change description

The **drni** keyword was changed to the **m-lag** keyword.

## Modified command: display port-security static-user connection

### Old syntax

```
display port-security static-user connection [[drni [local | peer]]
[interface interface-type interface-number | online-type
{ auth-fail-domain | critical-domain | preauth-domain | success } | slot
slot-number | user-name user-name] | { ip | ipv6 } ip-address | mac
mac-address]
```

### New syntax

```
display port-security static-user connection [[m-lag [local | peer]]
[interface interface-type interface-number | online-type
{ auth-fail-domain | critical-domain | preauth-domain | success } | slot
slot-number | user-name user-name] | { ip | ipv6 } ip-address | mac
mac-address]
```

### Views

Any view

### Change description

The **drni** keyword was changed to the **m-lag** keyword.

## Modified command: port-security drni load-sharing-mode

### Old syntax

```
port-security drni load-sharing-mode { centralized | distributed
{ even-mac | local | odd-mac } }

undo port-security drni load-sharing-mode
```

### New syntax

```
port-security m-lag load-sharing-mode { centralized | distributed
{ even-mac | local | odd-mac } }

undo port-security m-lag load-sharing-mode
```

### Views

Any view

### Change description

The **drni** keyword was changed to the **m-lag** keyword.

## Command changes in DHCP

## Modified command: display dhcp snooping drni-statistics

### Old syntax

```
display dhcp snooping drni-statistics [old-version]
```

### New syntax

```
display dhcp snooping m-lag-statistics [old-version]
```

### Views

Any view

### Change description

The **drni** keyword in this command was changed to the **m-lag** keyword.

Modified command: display dhcp snooping drni-status

### Old syntax

```
display dhcp snooping drni-status
```

### New syntax

```
display dhcp snooping m-lag-status
```

### Views

Any view

### Change description

The **drni** keyword in this command was changed to the **m-lag** keyword.

Modified command: reset dhcp snooping drni-statistics

### Old syntax

```
reset dhcp snooping drni-statistics
```

### New syntax

```
reset dhcp snooping m-lag-statistics
```

### Views

User view

### Change description

The **drni** keyword in this command was changed to the **m-lag** keyword.

## Command changes in DHCPv6

Modified command: display ipv6 dhcp snooping drni-statistics

### Old syntax

```
display ipv6 dhcp snooping drni-statistics [old-version]
```

### New syntax

```
display ipv6 dhcp snooping m-lag-statistics [old-version]
```

### Views

Any view

### Change description

The **drni** keyword in this command was changed to the **m-lag** keyword.

Modified command: display ipv6 dhcp snooping drni-status

### Old syntax

```
display ipv6 dhcp snooping drni-status
```

### New syntax

```
display ipv6 dhcp snooping m-lag-status
```

### Views

Any view

### Change description

The **drni** keyword in this command was changed to the **m-lag** keyword.

Modified command: reset ipv6 dhcp snooping drni-statistics

### Old syntax

```
reset ipv6 dhcp snooping drni-statistics
```

### New syntax

```
reset ipv6 dhcp snooping m-lag-statistics
```

### Views

User view

### Change description

The **drni** keyword in this command was changed to the **m-lag** keyword.

## Command changes in VXLAN

Modified command: display l2vpn drni arp

### Old syntax

```
display l2vpn drni arp [local | remote] [count | { public-instance |
vpn-instance vpn-instance-name } [count | ip-address]]
```

### New syntax

```
display l2vpn m-lag arp [local | remote] [count | { public-instance |
vpn-instance vpn-instance-name } [count | ip-address]]
```

### Views

Any view

Modified command: display l2vpn drni arp suppression

### Old syntax

```
display l2vpn drni arp suppression [local | remote] [count | vsi vsi-name
[count | ip-address]]
```

### New syntax

```
display l2vpn m-lag arp suppression [local | remote] [count | vsi
vsi-name [count | ip-address]]
```

### Views

Any view

Modified command: display l2vpn drni mac-address

#### Old syntax

```
display l2vpn drni mac-address [local | remote] [count | vsi vsi-name
[count | mac-address]]
```

#### New syntax

```
display l2vpn m-lag mac-address [local | remote] [count | vsi vsi-name
[count | mac-address]]
```

#### Views

Any view

Modified command: display l2vpn drni nd

#### Old syntax

```
display l2vpn drni nd [local | remote] [count | { public-instance |
vpn-instance vpn-instance-name } [ipv6-address | count]]
```

#### New syntax

```
display l2vpn m-lag nd [local | remote] [count | { public-instance |
vpn-instance vpn-instance-name } [ipv6-address | count]]
```

#### Views

Any view

Modified command: display l2vpn drni nd suppression

#### Old syntax

```
display l2vpn drni nd suppression [local | remote] [count | vsi vsi-name
[ipv6-address | count]]
```

#### New syntax

```
display l2vpn m-lag nd suppression [local | remote] [count | vsi vsi-name
[ipv6-address | count]]
```

#### Views

Any view

Modified command: display l2vpn service-instance

#### Syntax

```
display l2vpn service-instance [interface interface-type
interface-number [service-instance instance-id]] [verbose]
```

#### Views

Any view

#### Change description

A value for the **Type** field is modified.

Old value: **Dynamic (DRNI)**, representing a dynamic Ethernet service instance created on the IPL to match traffic by VLAN.

New value: **Dynamic (M-LAG)**, representing a dynamic Ethernet service instance created on the peer link to match traffic by VLAN.

Modified command: `display l2vpn vsi`

#### Syntax

```
display l2vpn vsi [name vsi-name] [verbose]
```

#### Views

Any view

#### Change description

A value for the **Type** field is modified.

Old value: **Dynamic (DRNI)**, representing a dynamic Ethernet service instance created on the IPL to match traffic by VLAN.

New value: **Dynamic (M-LAG)**, representing a dynamic Ethernet service instance created on the peer link to match traffic by VLAN.

## Command changes in EVPN

Modified command: `display evpn drni synchronized-mac`

#### Old syntax

```
display evpn drni synchronized-mac [vsi vsi-name] [count]
```

#### New syntax

```
display evpn m-lag synchronized-mac [vsi vsi-name] [count]
```

#### Views

Any view

Modified command: `evpn drni group`

#### Old syntax

```
evpn drni group { virtual-vtep-ipv4 / virtual-vtep-ipv6 }
```

#### New syntax

```
evpn m-lag group { virtual-vtep-ipv4 / virtual-vtep-ipv6 }
```

#### Views

System view

Modified command: `evpn drni local`

#### Old syntax

```
evpn drni local { local-ipv4-address remote remote-ipv4-address |
local-ipv6-address remote remote-ipv6-address }
```

#### New syntax

```
evpn m-lag local { local-ipv4-address remote remote-ipv4-address |
local-ipv6-address remote remote-ipv6-address }
```

## Views

System view

Modified command: l2vpn drni peer-link ac-match-rule vxlan-mapping

## Old syntax

```
l2vpn drni peer-link ac-match-rule vxlan-mapping
```

## New syntax

```
l2vpn m-lag peer-link ac-match-rule vxlan-mapping
```

## Views

System view

Modified command: drni local

## Old syntax

```
drni local local-ipv4-address remote remote-ipv4-address
```

## New syntax

```
m-lag local local-ipv4-address remote remote-ipv4-address
```

## Views

MVXLAN IPv4 address family view

Modified command: multicast-vpn vxlan drni local

## Old syntax

```
multicast-vpn vxlan drni local local-ipv4-address remote
remote-ipv4-address
```

## New syntax

```
multicast-vpn vxlan m-lag local local-ipv4-address remote
remote-ipv4-address
```

## Views

System view

Modified command: nexthop evpn-drni group-address

## Old syntax

```
nexthop evpn-drni group-address
```

## New syntax

```
nexthop evpn-m-lag group-address
```

## Views

BGP EVPN address family view

Modified command: source

#### Old syntax

```
source interface-type interface-number [evpn-drni-group]
```

#### New syntax

```
source interface-type interface-number [evpn-mlag-group]
```

#### Views

MVXLAN IPv4 address family view

## Modified feature: Using OSPF commands on interfaces

### Feature change description

As from this version, you can use interface-specific OSPF commands on VSI interfaces.

### Command changes

#### Syntax

```
ospf area
ospf authentication-mode
ospf bfd enable
ospf cost
ospf database-filter
ospf dr-priority
ospf fast-reroute lfa-backup
ospf mtu-enable
ospf network-type
ospf packet-size
ospf peer hold-max-cost duration
ospf peer sub-address enable
ospf prefix-suppression
ospf primary-path-detect bfd
ospf timer dead
ospf timer hello
ospf timer poll
ospf timer retransmit
ospf trans-delay
ospf ttl-security
```

#### Change description

These commands were added to VSI interface view.

# Modified feature: Configuring a static route

## Feature change description

As from this version, you do not need to specify a next hop address for a static route in the following situations:

- The static route belongs to the public network and a destination VPN instance is specified for the static route.
- The static route belongs to a VPN instance and a destination VPN instance or the public network is specified for the static route.

## Command changes

### Modified command: ip route-static

#### Old syntax

```
ip route-static dest-address { mask-length | mask } interface-type
interface-number [next-hop-address [nexthop-index index-string]]
[backup-interface interface-type interface-number [backup-nexthop
backup-nexthop-address] [permanent] | bfd { control-packet | echo-packet }
| permanent | track track-entry-number] [preference preference] [tag
tag-value] [recursive-lookup] [description text]
```

```
ip route-static dest-address { mask-length | mask } next-hop-address
[nexthop-index index-string] [recursive-lookup host-route] [bfd
control-packet bfd-source ip-address | permanent | track
track-entry-number] [preference preference] [tag tag-value]
[recursive-lookup] [description text]
```

```
ip route-static dest-address { mask-length | mask } vpn-instance
d-vpn-instance-name next-hop-address [nexthop-index index-string]
[recursive-lookup host-route] [bfd control-packet bfd-source ip-address
| permanent | track track-entry-number] [preference preference] [tag
tag-value] [recursive-lookup] [description text]
```

```
ip route-static group group-name interface-type interface-number
[next-hop-address] [backup-interface interface-type interface-number
[backup-nexthop backup-nexthop-address] [permanent] | bfd
{ control-packet | echo-packet } | permanent | track track-entry-number]
[preference preference] [tag tag-value] [recursive-lookup]
[description text]
```

```
ip route-static group group-name next-hop-address [recursive-lookup
host-route] [bfd control-packet bfd-source ip-address | permanent | track
track-entry-number] [preference preference] [tag tag-value]
[recursive-lookup] [description text]
```

```
ip route-static group group-name vpn-instance d-vpn-instance-name
next-hop-address [recursive-lookup host-route] [bfd control-packet
bfd-source ip-address | permanent | track track-entry-number]
[preference preference] [tag tag-value] [recursive-lookup]
[description text]
```

```
undo ip route-static { dest-address { mask-length | mask } | group
group-name } [interface-type interface-number [next-hop-address] |
```

```
next-hop-address | vpn-instance d-vpn-instance-name next-hop-address]
[preference preference]
```

## New syntax

```
ip route-static dest-address { mask-length | mask } interface-type
interface-number [next-hop-address [nexthop-index index-string]]
[backup-interface interface-type interface-number [backup-nexthop
backup-nexthop-address] [permanent] | bfd { control-packet | echo-packet }
| permanent | track track-entry-number] [preference preference] [tag
tag-value] [recursive-lookup] [description text]
```

```
ip route-static dest-address { mask-length | mask } next-hop-address
[nexthop-index index-string] [recursive-lookup host-route] [bfd
control-packet bfd-source ip-address | permanent | track
track-entry-number] [preference preference] [tag tag-value]
[recursive-lookup] [description text]
```

```
ip route-static dest-address { mask-length | mask } vpn-instance
d-vpn-instance-name next-hop-address [nexthop-index index-string]
[recursive-lookup host-route] [bfd control-packet bfd-source ip-address
| permanent | track track-entry-number] [preference preference] [tag
tag-value] [recursive-lookup] [description text]
```

```
ip route-static group group-name interface-type interface-number
[next-hop-address] [backup-interface interface-type interface-number
[backup-nexthop backup-nexthop-address] [permanent] | bfd
{ control-packet | echo-packet } | permanent | track track-entry-number]
[preference preference] [tag tag-value] [recursive-lookup]
[description text]
```

```
ip route-static group group-name next-hop-address [recursive-lookup
host-route] [bfd control-packet bfd-source ip-address | permanent | track
track-entry-number] [preference preference] [tag tag-value]
[recursive-lookup] [description text]
```

```
ip route-static group group-name vpn-instance d-vpn-instance-name
next-hop-address [recursive-lookup host-route] [bfd control-packet
bfd-source ip-address | permanent | track track-entry-number]
[preference preference] [tag tag-value] [recursive-lookup]
[description text]
```

```
undo ip route-static { dest-address { mask-length | mask } | group
group-name } [interface-type interface-number [next-hop-address] |
next-hop-address | vpn-instance d-vpn-instance-name next-hop-address]
[preference preference]
```

```
ip route-static { dest-address { mask-length | mask } | group group-name }
vpn-instance d-vpn-instance-name [preference preference] [tag tag-value]
[description text]
```

```
undo ip route-static { dest-address { mask-length | mask } | group
group-name } [vpn-instance d-vpn-instance-name] [preference
preference]
```

## Views

System view

## Change description

Before modification: If you specify a destination VPN instance, you must specify a next hop address. The device will search for the output interface in the destination VPN instance based on the configured next hop address for packets matching the static route.

After modification: If you specify a destination VPN instance, you do not need to specify a next hop address.

- If you do not specify a next hop address following the destination VPN instance, the device searches for the output interface in the specified destination VPN instance for packets matching the static route.
- If you specify a next hop address following the destination VPN instance, the device searches for the output interface in the specified destination VPN instance based on the specified next hop address for packets matching the static route.

## Modified command: ip route-static vpn-instance

### Old syntax

```
ip route-static vpn-instance s-vpn-instance-name dest-address
{ mask-length | mask } interface-type interface-number [next-hop-address
[nexthop-index index-string]] [backup-interface interface-type
interface-number [backup-nexthop backup-nexthop-address] [permanent] |
bfd { control-packet | echo-packet } | permanent | track track-entry-number]
[preference preference] [tag tag-value] [recursive-lookup]
[description text]
```

```
ip route-static vpn-instance s-vpn-instance-name dest-address
{ mask-length | mask } vpn-instance d-vpn-instance-name next-hop-address
[nexthop-index index-string] [recursive-lookup host-route] [bfd
control-packet bfd-source ip-address | permanent | track
track-entry-number] [preference preference] [tag tag-value]
[recursive-lookup] [description text]
```

```
ip route-static vpn-instance s-vpn-instance-name dest-address
{ mask-length | mask } next-hop-address [nexthop-index index-string]
[recursive-lookup host-route] [public] [bfd control-packet bfd-source
ip-address | permanent | track track-entry-number] [preference
preference] [tag tag-value] [recursive-lookup] [description text]
```

```
ip route-static vpn-instance s-vpn-instance-name group group-name
interface-type interface-number [next-hop-address] [bfd
{ control-packet | echo-packet } | backup-interface interface-type
interface-number [backup-nexthop backup-nexthop-address] [permanent]]
[preference preference] [tag tag-value] [recursive-lookup]
[description text]
```

```
ip route-static vpn-instance s-vpn-instance-name group group-name
next-hop-address [recursive-lookup host-route] [public] [bfd
control-packet bfd-source ip-address | permanent | track
track-entry-number] [preference preference] [tag tag-value]
[recursive-lookup] [description text]
```

```
ip route-static vpn-instance s-vpn-instance-name group group-name
vpn-instance d-vpn-instance-name next-hop-address [recursive-lookup
host-route] [bfd control-packet bfd-source ip-address | permanent | track
track-entry-number] [preference preference] [tag tag-value]
[recursive-lookup] [description text]
```

```
undo ip route-static vpn-instance s-vpn-instance-name { dest-address
{ mask-length | mask } | group group-name } [interface-type
interface-number [next-hop-address] | next-hop-address [public] |
vpn-instance d-vpn-instance-name next-hop-address] [preference
preference]
```

## New syntax

```
ip route-static vpn-instance s-vpn-instance-name dest-address
{ mask-length | mask } interface-type interface-number [next-hop-address
[nexthop-index index-string]] [backup-interface interface-type
interface-number [backup-nexthop backup-nexthop-address] [permanent] |
bfd { control-packet | echo-packet } | permanent | track track-entry-number]
[preference preference] [tag tag-value] [recursive-lookup]
[description text]

ip route-static vpn-instance s-vpn-instance-name dest-address
{ mask-length | mask } vpn-instance d-vpn-instance-name next-hop-address
[recursive-lookup host-route] [bfd { control-packet bfd-source
ip-address | static session-name } | permanent | track track-entry-number]
[preference preference] [tag tag-value] [recursive-lookup]
[description text]

ip route-static vpn-instance s-vpn-instance-name dest-address
{ mask-length | mask } next-hop-address [nexthop-index index-string]
[recursive-lookup host-route] [public] [bfd control-packet bfd-source
ip-address | permanent | track track-entry-number] [preference
preference] [tag tag-value] [recursive-lookup] [description text]

ip route-static vpn-instance s-vpn-instance-name group group-name
interface-type interface-number [next-hop-address] [bfd
{ control-packet | echo-packet } | backup-interface interface-type
interface-number [backup-nexthop backup-nexthop-address] [permanent]]
[preference preference] [tag tag-value] [recursive-lookup]
[description text]

ip route-static vpn-instance s-vpn-instance-name group group-name
next-hop-address [recursive-lookup host-route] [public] [bfd
control-packet bfd-source ip-address | permanent | track
track-entry-number] [preference preference] [tag tag-value]
[recursive-lookup] [description text]

ip route-static vpn-instance s-vpn-instance-name group group-name
vpn-instance d-vpn-instance-name next-hop-address [recursive-lookup
host-route] [bfd control-packet bfd-source ip-address | permanent | track
track-entry-number] [preference preference] [tag tag-value]
[recursive-lookup] [description text]

undo ip route-static vpn-instance s-vpn-instance-name { dest-address
{ mask-length | mask } | group group-name } [interface-type
interface-number [next-hop-address] | next-hop-address [public] |
vpn-instance d-vpn-instance-name next-hop-address] [preference
preference]

ip route-static vpn-instance s-vpn-instance-name { dest-address
{ mask-length | mask } | group group-name } { public | vpn-instance
d-vpn-instance-name } [preference preference] [tag tag-value]
[description text]

undo ip route-static vpn-instance s-vpn-instance-name { dest-address
{ mask-length | mask } | group group-name } [public | vpn-instance
d-vpn-instance-name] [preference preference]
```

## Views

System view

## Change description

Before modification: If you specify a destination VPN instance or the public network, you must specify a next hop address. The device will search for the output interface in the destination VPN instance or public network based on the configured next hop address for packets matching the static route.

After modification: If you specify a destination VPN instance or the public network, you do not need to specify a next hop address.

- If you do not specify a next hop address following the destination VPN instance, the device searches for the output interface in the specified destination VPN instance for packets matching the static route. Make sure the source and destination VPN instances are not the same one.
- If you specify a next hop address following the destination VPN instance, the device searches for the output interface in the specified destination VPN instance based on the specified next hop address for packets matching the static route.
- If you specify a next hop address with the **public** keyword, the next hop address is on the public network. The device searches for the output interface in the public network based on the next hop address for packets matching the static route.
- If you specify a next hop address without the **public** keyword, the device searches for the output interface in the source VPN instance based on the next hop address for packets matching the static route.
- If you specify the **public** keyword without specifying a next hop address, the device searches for the output interface in the public network for packets matching the static route.

## Modified feature: Configuring an IPv6 static route

### Feature change description

As from this version, you do not need to specify a next hop address for an IPv6 static route in the following situations:

- The IPv6 static route belongs to the public network and a destination VPN instance is specified for the IPv6 static route.
- The IPv6 static route belongs to a VPN instance and a destination VPN instance or the public network is specified for the IPv6 static route.

### Command changes

#### Modified command: `ipv6 route-static`

##### Old syntax

```
ipv6 route-static ipv6-address prefix-length { interface-type
interface-number [next-hop-address] [bfd { control-packet | echo-packet }
[bfd-source ipv6-address] | permanent | track track-entry-number] |
[vpn-instance d-vpn-instance-name] next-hop-address [recursive-lookup
host-route] [bfd control-packet bfd-source ipv6-address | permanent |
track track-entry-number] } [preference preference] [tag tag-value]
[recursive-lookup] [description text]

undo ipv6 route-static ipv6-address prefix-length [interface-type
interface-number [next-hop-address] | [vpn-instance
d-vpn-instance-name] next-hop-address] [preference preference]
```

```

ipv6 route-static vpn-instance s-vpn-instance-name ipv6-address
prefix-length { interface-type interface-number [next-hop-address] [bfd
{ control-packet | echo-packet } [bfd-source ipv6-address] | permanent |
track track-entry-number] | next-hop-address [public] [bfd
control-packet bfd-source ipv6-address | permanent | track
track-entry-number] | vpn-instance d-vpn-instance-name next-hop-address
[recursive-lookup host-route] [bfd control-packet bfd-source
ipv6-address | permanent | track track-entry-number] } [preference
preference] [tag tag-value] [recursive-lookup] [description text]

undo ipv6 route-static vpn-instance s-vpn-instance-name ipv6-address
prefix-length [interface-type interface-number [next-hop-address] |
next-hop-address [public] | vpn-instance d-vpn-instance-name
next-hop-address] [preference preference]

```

## New syntax

```

ipv6 route-static ipv6-address prefix-length interface-type
interface-number [next-hop-address] [bfd { control-packet | echo-packet }
[bfd-source ipv6-address] | permanent | track track-entry-number]
[preference preference] [tag tag-value] [recursive-lookup]
[description text]

ipv6 route-static ipv6-address prefix-length [vpn-instance
d-vpn-instance-name] next-hop-address [recursive-lookup host-route]
[bfd control-packet bfd-source ipv6-address | permanent | track
track-entry-number] [preference preference] [tag tag-value]
[recursive-lookup] [description text]

undo ipv6 route-static ipv6-address prefix-length [interface-type
interface-number [next-hop-address] | [vpn-instance
d-vpn-instance-name] next-hop-address] [preference preference]

ipv6 route-static ipv6-address prefix-length vpn-instance
d-vpn-instance-name [preference preference] [tag tag-value]
[description text]

undo ipv6 route-static ipv6-address prefix-length [vpn-instance
d-vpn-instance-name] [preference preference]

ipv6 route-static vpn-instance s-vpn-instance-name ipv6-address
prefix-length interface-type interface-number [next-hop-address] [bfd
{ control-packet | echo-packet } [bfd-source ipv6-address] | permanent |
track track-entry-number] } [preference preference] [tag tag-value]
[recursive-lookup] [description text]

ipv6 route-static vpn-instance s-vpn-instance-name ipv6-address
prefix-length { next-hop-address [recursive-lookup host-route] [public]
[bfd control-packet bfd-source ipv6-address | permanent | track
track-entry-number] | vpn-instance d-vpn-instance-name next-hop-address
[recursive-lookup host-route] [bfd control-packet bfd-source
ipv6-address | permanent | track track-entry-number] } [preference
preference] [tag tag-value] [recursive-lookup] [description text]

undo ipv6 route-static vpn-instance s-vpn-instance-name ipv6-address
prefix-length [interface-type interface-number [next-hop-address] |
next-hop-address [public] | vpn-instance d-vpn-instance-name
next-hop-address] [preference preference]

ipv6 route-static vpn-instance s-vpn-instance-name ipv6-address
prefix-length { public | vpn-instance d-vpn-instance-name } [preference
preference] [tag tag-value] [description text]

```

```
undo ipv6 route-static vpn-instance s-vpn-instance-name ipv6-address
prefix-length [public | vpn-instance d-vpn-instance-name] [preference
preference]
```

## Views

System view

## Change description

Before modification:

- When you configure an IPv6 static route for the public network, you must specify a next hop address if you specify a destination VPN instance. The device will search for the output interface in the destination VPN instance based on the configured next hop address for packets matching the IPv6 static route.
- When you configure an IPv6 static route for a VPN instance, you must specify a next hop address if you specify a destination VPN instance or the public network. The device will search for the output interface in the destination VPN instance or public network based on the configured next hop address for packets matching the IPv6 static route.

After modification:

- When you configure an IPv6 static route for the public network, you do not need to specify a next hop address if you specify a destination VPN instance.
  - If you do not specify a next hop address following the destination VPN instance, the device searches for the output interface in the specified destination VPN instance for packets matching the IPv6 static route.
  - If you specify a next hop address following the destination VPN instance, the device searches for the output interface in the specified destination VPN instance based on the specified next hop address for packets matching the IPv6 static route.
- When you configure an IPv6 static route for a VPN instance, you do not need to specify a next hop address if you specify a destination VPN instance or the public network.
  - If you do not specify a next hop address following the destination VPN instance, the device searches for the output interface in the specified destination VPN instance for packets matching the IPv6 static route. Make sure the source and destination VPN instances are not the same one.
  - If you specify a next hop address following the destination VPN instance, the device searches for the output interface in the specified destination VPN instance based on the specified next hop address for packets matching the IPv6 static route.
  - If you specify a next hop address with the **public** keyword, the next hop address is on the public network. The device searches for the output interface in the public network based on the next hop address for packets matching the IPv6 static route.
  - If you specify a next hop address without the **public** keyword, the device searches for the output interface in the source VPN instance based on the next hop address for packets matching the IPv6 static route.

If you specify the **public** keyword without specifying a next hop address, the device searches for the output interface in the public network for packets matching the IPv6 static route.

# ESS 6702

This release has the following changes:

- New feature: Specifying a security enhanced level
- New feature: Private VSI
- Modified feature: Configuring the global priority trust mode for VXLAN packets
- Modified feature: Configuring MAC authentication
- Modified feature: Disabling BGP from flushing all routes to the routing table
- Modified feature: Support for HMAC-SHA-256 OSPF/OSPFv3 authentication mode
- Modified feature: Displaying the hash keys used for link aggregation load sharing

## New feature: Specifying a security enhanced level

### Specifying a security enhanced level

#### About this task

The security enhanced level for the device can be 1 and 2, and level 2 indicates a higher security level. If the security enhanced level is set to 2, the following rules apply:

- SSL client policies and SSL server policies do not support cipher suites that contain DES, 3DES, MD5, RC4, and RC2.
- SSL client policies and SSL server policies do not support SSL protocol versions lower than TLS 1.1.
- SSL session renegotiation cannot be enabled.

You can specify a security enhanced level as needed.

#### Restrictions and guidelines

After you change the security enhanced level, for services associated with SSL policies, such as HTTP and SSL VPN, you must re-enable these services to update the associated policies.

#### Procedure

1. Enter system view.  
**system-view**
2. Specify the security enhanced level for the device.  
**security-enhanced level *level-value***  
By default, the security enhanced level is set to 2.

## Command reference

Use **security-enhanced level** to specify a security enhanced level for the device.

Use **undo security-enhanced level** to restore the default.

#### Syntax

**security-enhanced level *level-value***

```
undo security-enhanced level
```

## Default

The security enhanced level for the device is 2.

## Views

System view

## Predefined user roles

network-admin

## Parameters

*level-value*: Specifies the security enhanced level, which can be 1 and 2.

## Usage guidelines

The security enhanced level for the device can be 1 and 2, and level 2 indicates a higher security level. If the security enhanced level is set to 2, the following rules apply:

- SSL client policies and SSL server policies do not support cipher suites that contain DES, 3DES, MD5, RC4, and RC2.
- SSL client policies and SSL server policies do not support SSL protocol versions lower than TLS 1.1.
- SSL session renegotiation cannot be enabled.

You can specify a security enhanced level as needed.

After you change the security enhanced level, for services associated with SSL policies, such as HTTP and SSL VPN, you must re-enable these services to update the associated policies.

## Examples

```
Set the security enhanced level to 2 for the device.
```

```
<Sysname> system-view
```

```
[Sysname] security-enhanced level 2
```

# New feature: Private VSI

## Configuring private VSI

### About private VSI

In a data center network, typically VMs on a server are assigned to the same subnet. For security purposes, you can configure private VSI on the attached VTEP and private VLAN on the server to isolate the VMs based on VM attributes, operating system, and VM name. The server will send inter-VM traffic to the VTEP, and the VTEP will isolate traffic of different VMs or apply security policies as configured.

The private VSI feature uses a two-tier VSI structure, as shown in [Figure 1](#). Each primary VSI is associated with one or multiple secondary VSIs, and the secondary VSIs communicate at Layer 3 through the primary VSI. A secondary VSI can be one of the following types:

- **Community VSI**—A secondary VSI whose ACs have Layer 2 connectivity to one another.
- **Isolated VSI**—A secondary VSI whose ACs are isolated at Layer 2.

**Figure 1 Private VSI example**

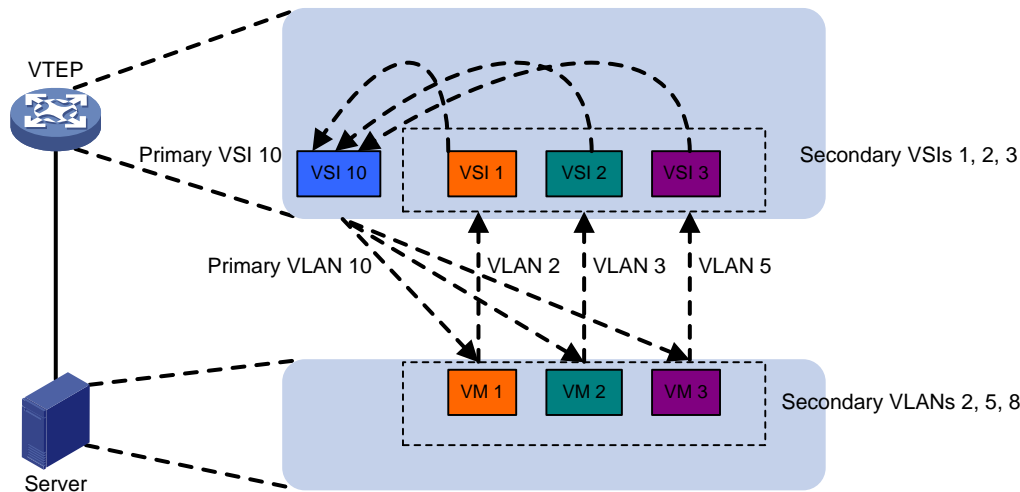


Figure 1 shows a private VSI application example. The private VLAN feature is configured on the server as follows:

- VLAN 10 is a primary VLAN.
- VLAN 2, VLAN 5, and VLAN 8 are secondary VLANs of VLAN 10.

The following VSIs are configured on the VTEP:

| VSI    | VSI type      | VXLAN ID | Mapped VLAN |
|--------|---------------|----------|-------------|
| VSI 10 | Primary VSI   | 10       | VLAN 10     |
| VSI 2  | Community VSI | 2        | VLAN 2      |
| VSI 5  | Community VSI | 5        | VLAN 5      |
| VSI 8  | Community VSI | 8        | VLAN 8      |

To enable Layer 3 communication among the secondary VSIs, VSI-interface 1 is configured as the gateway interface for all VSIs, and local ARP and ND proxy is enabled on VSI-interface 1. The VTEP sends all packets destined for the VMs on the server only through the ACs of the primary VSI.

The following uses packet forwarding from VM 1 to VM 3 as an example to introduce the inter-VM communication procedure:

1. VM 1 sends a packet in VLAN 2. The destination IP address of the packet is the IP address of VM 3, and the destination MAC address is the MAC address of VSI-interface 1.
2. The server forwards the packet to the VTEP.
3. The VTEP receives the packet from an AC of VSI 1 and forwards the packet to VSI-interface 1.
4. VSI-interface 1 looks up the ARP table for the outgoing interface, encapsulates the packet, and forwards it through an AC of VSI 10 on the outgoing interface.
5. The server forwards the packet to VM 3.

## Restrictions and guidelines

You can configure only one-to-one mappings between private VLAN sets and private VSI sets.

As a best practice, configure Ethernet service instances for VSIs as follows:

- Configure a primary VSI to match the packets of a primary VLAN on the peer server.

- Configure a secondary VSI of the primary VSI to match the packets of a secondary VLAN of the primary VLAN on the peer server.

To ensure correct Layer 3 forwarding, you must use the same VSI interface as the gateway interface for a primary VSI and its associated secondary VSIs. You must configure the VSI interface and the primary VSI as follows:

- Enable local ARP or ND proxy on the VSI interface.
- Execute the **gateway subnet** command to assign the primary VSI to the subnet where the gateway address belongs.

When you map ACs to VSIs, follow these restrictions and guidelines:

- On a physical interface, you can map only one AC to a primary VSI.
- If an AC on a physical interface has been mapped to a secondary VSI, you must map another AC on that physical interface to the associated primary VSI. If you fail to do so, private VSI configuration does not take effect.

You cannot associate a secondary VSI with multiple primary VSIs. You can associate a maximum of eight secondary VSIs with a primary VSI.

## Configuring a primary VSI

1. Enter system view.  
**system-view**
2. Enter VSI view.  
**vsi** *vsi-name*
3. Configure the VSI as a primary VSI and enable the private VSI feature.  
**private-vsi primary**  
By default, the private VSI feature is disabled.
4. Specify secondary VSIs.  
**private-vsi secondary** *vsi-name*&<1-8>  
By default, a primary VSI is not associated with any secondary VSIs.

## Configuring a secondary VSI

1. Enter system view.  
**system-view**
2. Enter VSI view.  
**vsi** *vsi-name*
3. Configure the VSI as a secondary VSI and enable the private VSI feature.  
**private-vsi { community | isolated }**  
By default, the private VSI feature is disabled.

## Command reference

### private-vsi

Use **private-vsi** to configure a VSI as a primary or secondary VSI and enable the private VSI feature.

Use **undo private-vsi** to disable the private VSI feature.

## Syntax

```
private-vsi { community | isolated | primary }
undo private-vsi
```

## Default

The private VSI feature is disabled.

## Views

VSI view

## Predefined user roles

network-admin

## Parameters

**community**: Specifies the VSI as a secondary VSI whose ACs have Layer 2 connectivity to one another.

**isolated**: Specifies the VSI as a secondary VSI whose ACs are isolated at Layer 2.

**primary**: Specifies the VSI as a primary VSI.

## Usage guidelines

In a data center network, typically VMs on a server are assigned to the same subnet. For security purposes, you can configure private VSI on the attached VTEP and private VLAN on the server to isolate the VMs based on VM attributes, operating system, and VM name. The server will send inter-VM traffic to the VTEP, and the VTEP will isolate traffic of different VMs or apply security policies as configured.

## Examples

# Configure VSI **vpn1** as a primary VSI, and enable the private VSI feature.

```
<Sysname> system-view
[Sysname] vsi vpn1
[Sysname-vsi-vpn1] private-vsi primary
```

## private-vsi secondary

Use **private-vsi secondary** to associate a secondary VSI with a primary VSI.

Use **undo private-vsi secondary** to remove a secondary VSI from a primary VSI.

## Syntax

```
private-vsi secondary vsi-name&<1-8>
undo private-vsi secondary vsi-name&<1-8>
```

## Default

A primary VSI is not associated with any secondary VSIs.

## Views

VSI view

## Predefined user roles

network-admin

## Parameters

**vsi-name**: Specifies a space-separated list of up to eight VSI names. Each VSI name is a case-sensitive string of 1 to 31 characters.

## Usage guidelines

Each primary VSI is associated with one or multiple secondary VSIs, and the secondary VSIs communicate at Layer 3 through the primary VSI.

## Examples

# Associate secondary VSI **vpn2** with primary VSI **vpn1**.

```
<Sysname> system-view
```

```
[Sysname] vsi vpn1
```

```
[Sysname-vsi-vpn1] private-vsi primary
```

```
[Sysname-vsi-vpn1] private-vsi secondary vpn2
```

# Modified feature: Configuring the global priority trust mode for VXLAN packets

## Feature change description

As from this version, you can configure the switch to trust the DSCP priority in the outer IP header of VXLAN packets.

## Command changes

### Modified command: qos trust (system view)

#### Old syntax

```
qos trust tunnel-dot1p
```

```
undo qos trust tunnel-dot1p
```

#### New syntax

```
qos trust { tunnel-dot1p | tunnel-dscp }
```

```
undo qos trust { tunnel-dot1p | tunnel-dscp }
```

#### Change description

Before modification: The **tunnel-dscp** keyword is not supported.

After modification: The **tunnel-dscp** keyword is supported.

**tunnel-dscp**: Trusts the DSCP value in the outer IP header of VXLAN packets for priority mapping.

# Modified feature: Configuring MAC authentication

## Feature change description

As from this version, when local QoS ID settings exist, make sure the local QoS ID settings meet the following requirements:

- The local QoS ID specified by using the **if-match qos-local-id** command is smaller than 3000.
- The local QoS ID specified by using the **remark qos-local-id** command is smaller than 3000.

## Command changes

### Syntax

```
mac-authentication
remark qos-local-id [egress-active] local-id-value
if-match qos-local-id local-id-value [qppb-manipulation]
```

### Change description

Before modification: MAC authentication does not have configuration conflicts with the **if-match qos-local-id** or **remark qos-local-id** command when MAC authentication users come online.

After modification: When local QoS ID settings exist, make sure the local QoS ID settings meet the following requirements:

- The local QoS ID specified by using the **if-match qos-local-id** command is smaller than 3000.
- The local QoS ID specified by using the **remark qos-local-id** command is smaller than 3000.

## Modified feature: Disabling BGP from flushing all routes to the routing table

### Feature change description

As from this release, the **routing-table bgp-rib-only** command can disable BGP from flushing all routes to the routing table, including redistributed routes and routes received from peers and peer groups.

## Command changes

### Modified command: routing-table bgp-rib-only

#### Old syntax

```
routing-table bgp-rib-only [route-policy route-policy-name]
```

#### New syntax

```
routing-table bgp-rib-only [all] [route-policy route-policy-name]
```

#### Views

BGP IPv4 unicast address family view

BGP IPv6 unicast address family view

BGP-VPN IPv4 unicast address family view

BGP-VPN IPv6 unicast address family view

### Change description

Before modification: The **routing-table bgp-rib-only** command can only disable BGP from flushing routes received from peers and peer groups to the routing table.

After modification: The **routing-table bgp-rib-only** command supports the **all** keyword. This keyword disables BGP from flushing all routes to the routing table, including redistributed routes

and routes received from peers and peer groups. If you do not specify this keyword, BGP does not flush the routes received from the specified peer or peer group to the routing table.

## Modified feature: Support for HMAC-SHA-256 OSPF/OSPFv3 authentication mode

### Feature change description

As from this release, HMAC-SHA-256 OSPF/OSPFv3 authentication mode is supported.

### Command changes

#### Modified command: authentication-mode

##### Old syntax

For MD5/HMAC-MD5 authentication:

```
authentication-mode { hmac-md5 | md5 } [key-id { cipher | plain } string]
undo authentication-mode [{ hmac-md5 | md5 } [key-id]]
```

##### New syntax

For MD5/HMAC-MD5/HMAC-SHA-256 authentication:

```
authentication-mode { hmac-md5 | hmac-sha-256 | md5 } [key-id { cipher |
plain } string]
undo authentication-mode [{ hmac-md5 | hmac-sha-256 | md5 } [key-id]]
```

##### Views

OSPF area view

##### Change description

Before modification: This command does not support the **hmac-sha-256** keyword. You cannot enable the HMAC-SHA-256 OSPF authentication mode.

After modification: This command supports the **hmac-sha-256** keyword. You can enable the HMAC-SHA-256 OSPF authentication mode by specifying the **hmac-sha-256** keyword.

#### Modified command: ospf authentication-mode

##### Old syntax

For MD5/HMAC-MD5 authentication:

```
ospf authentication-mode { hmac-md5 | md5 } [key-id { cipher | plain }
string]
undo ospf authentication-mode [{ hmac-md5 | md5 } [key-id]]
```

##### New syntax

For MD5/HMAC-MD5/HMAC-SHA-256 authentication:

```
ospf authentication-mode { hmac-md5 | hmac-sha-256 | md5 } [key-id { cipher
| plain } string]
```

```
undo ospf authentication-mode [{ hmac-md5 | hmac-sha-256 | md5 }
[key-id]]
```

## Views

Interface view

## Change description

Before modification: This command does not support the **hmac-sha-256** keyword. You cannot enable the HMAC-SHA-256 OSPF authentication mode.

After modification: This command supports the **hmac-sha-256** keyword. You can enable the HMAC-SHA-256 OSPF authentication mode by specifying the **hmac-sha-256** keyword.

## Modified command: vlink-peer

### Old syntax

```
vlink-peer router-id [dead seconds | hello seconds | [authentication-none
/ { hmac-md5 | md5 } [key-id { cipher | plain } string] | keychain
keychain-name | simple [{ cipher | plain } string]] | retransmit seconds
| trans-delay seconds] *

undo vlink-peer router-id [dead | hello | [authentication-none /
{ hmac-md5 | md5 } [key-id] | keychain] | retransmit | simple | trans-delay
*

```

### New syntax

```
vlink-peer router-id [dead seconds | hello seconds | [authentication-none
/ { hmac-md5 | hmac-sha-256 | md5 } [key-id { cipher | plain } string] |
keychain keychain-name | simple [{ cipher | plain } string]] | retransmit
seconds | trans-delay seconds] *

undo vlink-peer router-id [dead | hello | [authentication-none /
{ hmac-md5 | hmac-sha-256 | md5 } [key-id] | keychain] | retransmit |
simple | trans-delay] *

```

## Views

OSPF area view

## Change description

Before modification: This command does not support the **hmac-sha-256** keyword. You cannot enable the HMAC-SHA-256 OSPF authentication mode.

After modification: This command supports the **hmac-sha-256** keyword. You can enable the HMAC-SHA-256 OSPF authentication mode by specifying the **hmac-sha-256** keyword.

## Modified command: authentication-mode

### Old syntax

```
authentication-mode keychain keychain-name
undo authentication-mode
```

### New syntax

For HMAC-SHA-256 authentication:

```
authentication-mode hmac-sha-256 key-id { cipher | plain } string
undo authentication-mode
```

For keychain authentication:

```
authentication-mode keychain keychain-name
undo authentication-mode
```

## Views

OSPFv3 area view

## Change description

Before modification: This command does not support the **hmac-sha-256** keyword. You cannot enable the HMAC-SHA-256 OSPFv3 authentication mode.

After modification: This command supports the **hmac-sha-256** keyword. You can enable the HMAC-SHA-256 OSPFv3 authentication mode by specifying the **hmac-sha-256** keyword.

## Modified command: ospfv3 authentication-mode

### Old syntax

```
ospfv3 authentication-mode keychain keychain-name [instance
instance-id]
undo ospfv3 authentication-mode [instance instance-id]
```

### New syntax

For HMAC-SHA-256 authentication:

```
ospfv3 authentication-mode hmac-sha-256 key-id { cipher | plain } string
[instance instance-id]
undo ospfv3 authentication-mode [instance instance-id]
```

For keychain authentication:

```
ospfv3 authentication-mode keychain keychain-name [instance
instance-id]
undo ospfv3 authentication-mode [instance instance-id]
```

## Views

Interface view

## Change description

Before modification: This command does not support the **hmac-sha-256** keyword. You cannot enable the HMAC-SHA-256 OSPFv3 authentication mode.

After modification: This command supports the **hmac-sha-256** keyword. You can enable the HMAC-SHA-256 OSPFv3 authentication mode by specifying the **hmac-sha-256** keyword.

## Modified command: vlink-peer

### Old syntax

```
vlink-peer router-id [dead seconds | hello seconds | instance instance-id
| ipsec-profile profile-name | keychain keychain-name | retransmit seconds
| trans-delay seconds] *
undo vlink-peer router-id [dead | hello | ipsec-profile | keychain |
retransmit | trans-delay] *
```

## New syntax

```
vlink-peer router-id [dead seconds | hello seconds | instance instance-id | ipsec-profile profile-name | [hmac-sha-256 key-id { cipher | plain } string | keychain keychain-name] | retransmit seconds | trans-delay seconds] *

undo vlink-peer router-id [dead | hello | ipsec-profile | [hmac-sha-256 | keychain] | retransmit | trans-delay] *
```

## Views

OSPFv3 area view

## Change description

Before modification: This command does not support the **hmac-sha-256** keyword. You cannot enable the HMAC-SHA-256 OSPFv3 authentication mode.

After modification: This command supports the **hmac-sha-256** keyword. You can enable the HMAC-SHA-256 OSPFv3 authentication mode by specifying the **hmac-sha-256** keyword.

## Modified command: sham-link (OSPF area view)

## Old syntax

```
sham-link source-ip-address destination-ip-address [cost cost-value | dead dead-interval | hello hello-interval | { authentication-none | { hmac-md5 | md5 } [key-id { cipher | plain } string] | keychain keychain-name | retransmit retrans-interval | simple [{ cipher | plain } string] } | trans-delay delay | ttl-security hops hop-count] *

undo sham-link source-ip-address destination-ip-address [cost | dead | hello | { authentication-none | { hmac-md5 | md5 } [key-id] | keychain | simple } | retransmit | trans-delay | ttl-security] *
```

## New syntax

```
sham-link source-ip-address destination-ip-address [cost cost-value | dead dead-interval | hello hello-interval | { authentication-none | { hmac-md5 | hmac-sha-256 | md5 } [key-id { cipher | plain } string] | keychain keychain-name | simple [{ cipher | plain } string] } | retransmit retrans-interval | trans-delay delay | ttl-security hops hop-count] *

undo sham-link source-ip-address destination-ip-address [cost | dead | hello | { authentication-none | { hmac-md5 | hmac-sha-256 | md5 } [key-id] | keychain | simple } | retransmit | trans-delay | ttl-security] *
```

## Views

OSPF area view

## Change description

Before modification: This command does not support the **hmac-sha-256** keyword. You cannot enable the HMAC-SHA-256 OSPF authentication mode.

After modification: This command supports the **hmac-sha-256** keyword. You can enable the HMAC-SHA-256 OSPF authentication mode by specifying the **hmac-sha-256** keyword.

## Modified command: sham-link (OSPFv3 area view)

### Old syntax

```
sham-link source-ipv6-address destination-ipv6-address [cost cost-value
| dead dead-interval | hello hello-interval | instance instance-id |
ipsec-profile profile-name | keychain keychain-name | retransmit
retrans-interval | trans-delay delay] *
```

```
undo sham-link source-ipv6-address destination-ipv6-address [cost | dead
| hello | ipsec-profile | keychain | retransmit | trans-delay] *
```

### New syntax

```
sham-link source-ipv6-address destination-ipv6-address [cost cost-value
| dead dead-interval | hello hello-interval | instance instance-id |
ipsec-profile profile-name | { hmac-sha-256 key-id { cipher | plain }
string | keychain keychain-name } | retransmit retrans-interval |
trans-delay delay] *
```

```
undo sham-link source-ipv6-address destination-ipv6-address [cost | dead
| hello | ipsec-profile | { hmac-sha-256 | keychain } | retransmit |
trans-delay] *
```

### Views

OSPFv3 area view

### Change description

Before modification: This command does not support the **hmac-sha-256** keyword. You cannot enable the HMAC-SHA-256 OSPFv3 authentication mode.

After modification: This command supports the **hmac-sha-256** keyword. You can enable the HMAC-SHA-256 OSPFv3 authentication mode by specifying the **hmac-sha-256** keyword.

## Modified feature: Displaying the hash keys used for link aggregation load sharing

### Feature change description

From this software version on, the device displays the hash keys used for load sharing Layer 2 and Layer 3 traffic over aggregate links.

### Command changes

## Modified command: display link-aggregation load-sharing mode

### Syntax

```
display link-aggregation load-sharing mode [interface
[{ bridge-aggregation | route-aggregation } interface-number]]
```

### Views

Any view

## Change description

Before modification: The **display link-aggregation load-sharing mode** command does not display load sharing hash keys for Layer 2 or Layer 3 traffic.

```
<Sysname> display link-aggregation load-sharing mode
MAC-in-MAC traffic load-sharing mode:
Outer (default)
Link-aggregation load-sharing algorithm:
5 (default)
Link-aggregation load-sharing offset:
0 (default)
Link-aggregation load-sharing seed:
0x0 (default)
Tunneled traffic load-sharing mode:
Outer (default)
Link-aggregation load-sharing mode:
 Layer 2 traffic: packet type-based sharing
 Layer 3 traffic: packet type-based sharing
```

After modification: The **display link-aggregation load-sharing mode** command displays load sharing hash keys for Layer 2 and Layer 3 traffic.

```
<Sysname> display link-aggregation load-sharing mode
MAC-in-MAC traffic load-sharing mode:
Outer (default)
Link-aggregation load-sharing algorithm:
5 (default)
Link-aggregation load-sharing offset:
0 (default)
Link-aggregation load-sharing seed:
0x0 (default)
Tunneled traffic load-sharing mode:
Outer (default)
Link-aggregation load-sharing mode:
 Layer 2 traffic: destination-mac address, source-mac address
 ethernet-type
 Layer 3 traffic: destination-ip address, source-ip address
 destination-port, source-port
 ip-protocol
```

# Release 2717

This release has the following changes:

- New feature: Interface alarm functions
- New feature: Establishing neighbors through the secondary IP address of an interface
- New feature: Advertising only the global unicast address in the NEXT\_HOP attribute
- New feature: Assigning a VXLAN tunnel interface to a VXLAN tunnel group
- New feature: Setting VLAN tags for sent protocol packets on an interface
- New feature: Configuring the system ID
- New feature: IPv6 duplicate detection on duplicate addresses
- New feature: BGP route re-origination
- New feature: Sending RA messages over VXLAN tunnels
- New feature: SAVA
- New feature: IPv6 IS-IS link tag
- New feature: OSPFv3 link tag
- New feature: Displaying a message on enablement of the FTP or Telnet server
- New feature: Enabling OSPF to advertise the maximum link cost to neighbors
- New feature: Enabling IS-IS to advertise the maximum link cost to neighbors
- New feature: Enabling OSPFv3 to advertise the maximum link cost to neighbors
- New feature: Enabling BGP to advertise a default route to a peer or peer group
- New feature: SAVI
- New feature: Setting the maximum number of IPv6SG bindings on an interface
- New feature: Setting the interval at which the SNMP module examines the system configuration for changes
- New feature: Displaying SNMP notifications drop records
- New feature: Displaying SNMP notifications sending records
- New feature: Clearing all records from the SNMP trap buffer
- New feature: PWDCTL messages
- New feature: LIPC messages
- New feature: Applying an SSL server policy to the NETCONF over SOAP over HTTPS service
- New feature: Policing ND protocol packet to the control plane
- New feature: RA message transmission over VXLAN tunnels between distributed EVPN gateways
- New feature: MAC move logging for VSIs
- New feature: MAC learning limits for VSIs
- New feature: DRNI in EVPN-DCI
- New feature: L2VPN statistics collection interval
- New feature: Unsolicited NA learning
- New feature: IPv6 enhancement for OpenFlow flow entries
- New feature: IPv6-matching QoS policy
- New feature: Enabling BGP to not flush specific routes to the routing table

- New feature: Enabling BGP to reset peer sessions gracefully
- New feature: Setting the extra time to wait after the restart timer expires
- New feature: Enabling mandatory weak password change
- New feature: Configuring ND packet rate limit
- New feature: gRPC logging in dial-in mode
- New feature: gRPC logging in dial-out mode
- New feature: Creating frame match criteria based on VXLAN IDs for dynamic ACs on the IPL of an EVPN DR system
- New feature: Removing Layer 3 VPN route targets from BGP EVPN routes for VXLAN mapping
- New feature: Enabling asymmetric IRB mode for EVPN VXLAN traffic forwarding
- New feature: Setting the Gateway IP Address field for EVPN IP prefix advertisement routes
- New feature: Enabling BGP to redistribute all IGP routes with the same destination address
- New feature: Policing ND protocol packets
- New feature: Using an ACL to control TCP connections from IPv4 HTTP or HTTPS clients
- New feature: Enabling Layer 3 packet statistics counting
- New feature: Specifying the PKI domain for secure communication with collectors
- New feature: DRNI MAD
- New feature: Configuring basic DHCPv6 snooping features on a VXLAN network
- New feature: Configuring BGP VPNv6 routes to use private network next hops
- New feature: Configuring BGP route replication between public and VPN instances
- New feature: Displaying IPv6 multicast group information that MLD snooping learns from EVPN
- New feature: Enabling password control blacklist for all user line types
- New feature: Configuring ND attack detection for a VSI
- New feature: MVXLAN
- Modified feature: Configuring LLDP to advertise IPv4 and IPv6 management addresses at the same time
- Modified feature: Specifying a padding mode for the Interface-ID option on the DHCPv6 relay agent
- Modified feature: Displaying monitor link group information
- Modified feature: Displaying the path that the IPv6 packets traverse from source to destination
- Modified feature: Exchanging both IPv4 and IPv6 routes with a BGP peer
- Modified feature: Enabling load balancing for BGP routes that have different AS\_PATH attributes of the same length
- Modified feature: Enabling BGP to exchange labeled routes with a peer or peer group
- Modified feature: Configuring a static neighbor entry
- Modified feature: Modifying information in received BGP EVPN IP prefix advertisement routes on an ED
- Modified feature: Configuring passive TCP connection settings for a VXLAN OVSDb VTEP
- Modified feature: Setting the password for local password authentication
- Modified feature: Enabling an ED to replace the L3 VXLAN ID, RD, and route targets of BGP EVPN routes
- Modified feature: Displaying Ethernet service instance information
- Modified feature: Configuring OSPF FRR
- Modified feature: Displaying routing table information

- Modified feature: Configuring the maximum number of Add-Path optimal routes that can be advertised to all peers
- Modified feature: Configuring the BGP Additional Paths capabilities
- Modified feature: Configuring the maximum number of Add-Path optimal routes that can be advertised to a peer or peer group
- Modified feature: Password control restriction features
- Modified feature: Flow-mirroring traffic to interfaces
- Modified feature: Configuring ARP scanning
- Modified feature: NETCONF logging
- Modified feature: Enabling route router MAC replacement for a peer or peer group on an EVPN-DCI ED
- Modified feature: Value range change for the extended community attribute list number in an RR reflection policy
- Modified feature: Allowing a DR interface to be up when it is the only member of a DR group
- Modified feature: Configuring a static route
- Modified feature: Configuring an IPv6 static route
- Modified feature: Displaying OSPFv3 LSDB information
- Modified feature: Configuring BGP extended community lists
- Modified feature: Replicating VLINK direct routes
- Modified feature: Specifying a controller for an OpenFlow switch and configuring the main connection to the controller
- Modified feature: Displaying the CPU usage statistics
- Modified feature: Displaying memory usage information
- Modified feature: Password complexity checking
- Modified feature: Configuring kernel thread deadlock detection
- Modified feature: Configuring a traffic redirecting action in a QoS traffic behavior
- Modified feature: DRNI IPP configuration

## New feature: Interface alarm functions

### Configuring interface alarm functions

#### About this task

With the interface alarm functions enabled, when the number of error packets on an interface in normal state within the specified interval exceeds the upper threshold, the interface generates an upper threshold exceeding alarm and enters the alarm state. When the number of error packets on an interface in the alarm state within the specified interval drops below the lower threshold, the interface generates a recovery alarm and restores to the normal state.

#### Restrictions and guidelines

You can configure the interface alarm parameters in system view and interface view.

- The configuration in system view takes effect on all interfaces of the specified slot. The configuration in interface view takes effect only on the current interface.
- For an interface, the configuration in interface view takes priority, and the configuration in system view is used only when no configuration is made in interface view.

An interface that is shut down because of error packet alarms cannot automatically recover. To bring up the interface, execute the **undo shutdown** command on the interface.

## Enabling interface alarm functions

1. Enter system view.  
**system-view**
2. Enable alarm functions for the interface monitoring module.  
**snmp-agent trap enable ifmonitor [ crc-error | input-error | input-usage | output-error | output-usage ] \***  
By default, all alarm functions are enabled for interfaces.

## Configuring CRC error packet parameters

1. Enter system view.  
**system-view**
2. Configure global CRC error packet alarm parameters.  
**ifmonitor crc-error slot slot-number high-threshold high-value low-threshold low-value interval interval [ shutdown ]**  
By default, the upper threshold is 1000, the lower threshold is 100, and the statistics collection and comparison interval is 10 seconds for CRC error packets.
3. Enter Ethernet interface view.  
**interface interface-type interface-number**
4. Configure CRC error packet alarm parameters for the interface.  
**port ifmonitor crc-error [ ratio ] high-threshold high-value low-threshold low-value interval interval [ shutdown ]**  
By default, an interface uses the global CRC error packet alarm parameters.

## Configuring input error packet alarm parameters

1. Enter system view.  
**system-view**
2. Configure global input error packet alarm parameters.  
**ifmonitor input-error slot slot-number high-threshold high-value low-threshold low-value interval interval [ shutdown ]**  
By default, the upper threshold is 1000, the lower threshold is 100, and the statistics collection and comparison interval is 10 seconds for input error packets.
3. Enter Ethernet interface view.  
**interface interface-type interface-number**
4. Configure input error packet alarm parameters for the interface.  
**port ifmonitor input-error high-threshold high-value low-threshold low-value interval interval [ shutdown ]**  
By default, an interface uses the global input error packet alarm parameters.

## Configuring output error packet alarm parameters

1. Enter system view.  
**system-view**
2. Configure global output error packet alarm parameters.  
**ifmonitor output-error slot slot-number high-threshold high-value low-threshold low-value interval interval [ shutdown ]**

By default, the upper threshold is 1000, the lower threshold is 100, and the statistics collection and comparison interval is 10 seconds for output error packets.

3. Enter Ethernet interface view.

```
interface interface-type interface-number
```

4. Configure output error packet alarm parameters.

```
port ifmonitor output-error high-threshold high-value low-threshold low-value interval interval [shutdown]
```

By default, an interface uses the global output error packet alarm parameters.

## Configuring input bandwidth usage alarm parameters

1. Enter system view.

```
system-view
```

2. Configure global input bandwidth usage alarm parameters.

```
ifmonitor input-usage slot slot-number high-threshold high-value low-threshold low-value
```

By default, the upper threshold is 90 and the lower threshold is 80 for input bandwidth usage alarms.

3. Enter Ethernet interface view.

```
interface interface-type interface-number
```

4. Configure input bandwidth usage alarm parameters.

```
port ifmonitor input-usage high-threshold high-value low-threshold low-value
```

By default, an interface uses the global input bandwidth usage alarm parameters.

## Configuring output bandwidth usage alarm parameters

1. Enter system view.

```
system-view
```

2. Configure global output bandwidth usage alarm parameters.

```
ifmonitor output-usage slot slot-number high-threshold high-value low-threshold low-value
```

By default, the upper threshold is 90 and the lower threshold is 80 for output bandwidth usage alarms.

3. Enter Ethernet interface view.

```
interface interface-type interface-number
```

4. Configure output bandwidth usage alarm parameters.

```
port ifmonitor output-usage high-threshold high-value low-threshold low-value
```

By default, an interface uses the global output bandwidth usage alarm parameters.

## Command reference

### ifmonitor crc-error

Use **ifmonitor crc-error** to configure global CRC error packet alarm parameters.

Use **undo ifmonitor crc-error** to restore the default.

## Syntax

```
ifmonitor crc-error slot slot-number high-threshold high-value
low-threshold low-value interval interval [shutdown]

undo ifmonitor crc-error slot slot-number
```

## Default

The upper threshold is 1000, the lower threshold is 100, and the statistics collection and comparison interval is 10 seconds for CRC error packet alarms.

## Views

System view

## Predefined user roles

network-admin

## Parameters

**high-threshold** *high-value*: Specifies the upper threshold for CRC error packet alarms, in the range of 1 to 4294967295 packets.

**low-threshold** *low-value*: Specifies the lower threshold for CRC error packet alarms, in the range of 1 to 4294967295 packets.

**interval** *interval*: Specifies the statistics collection and comparison interval for CRC error packets, in the range of 1 to 65535 seconds.

**shutdown**: Shuts down an interface when the number of incoming CRC error packets on the interface exceeds the upper threshold. Then, the interface stops forwarding all packets. To recover the interface, execute the **undo shutdown** command on the interface. If you do not specify this keyword, an upper threshold exceeding alarm is generated and the interface enters the alarm state when the number of incoming CRC error packets exceeds the upper threshold on the interface.

**slot** *slot-number*: Specifies an IRF member device by its member ID.

## Usage guidelines

With the CRC error packet alarm function enabled, when the number of incoming CRC error packets on an interface in normal state within the specified interval exceeds the upper threshold, the interface generates an upper threshold exceeding alarm and enters the alarm state. When the number of incoming CRC error packets on an interface in the alarm state within the specified interval drops below the lower threshold, the interface generates a recovery alarm and restores to the normal state.

You can configure the CRC error packet alarm parameters in system view and interface view.

- The configuration in system view takes effect on all interfaces of the specified slot. The configuration in interface view takes effect only on the current interface.
- For an interface, the configuration in interface view takes priority, and the configuration in system view is used only when no configuration is made in interface view.

When you execute this command multiple times, the most recent configuration takes effect.

## Examples

# Set the upper threshold to 5000, lower threshold to 400, and statistics collection and comparison interval to 6 seconds for CRC error packet alarms.

```
<Sysname> system-view
```

```
[Sysname] ifmonitor crc-error slot 1 high-threshold 5000 low-threshold 400 interval 6
```

## Related commands

```
snmp-agent trap enable ifmonitor
```

## ifmonitor input-error

Use **ifmonitor input-error** to configure global input error packet alarm parameters.

Use **undo ifmonitor input-error** to restore the default.

### Syntax

```
ifmonitor input-error slot slot-number high-threshold high-value
low-threshold low-value interval interval [shutdown]
```

```
undo ifmonitor input-error slot slot-number
```

### Default

The upper threshold is 1000, the lower threshold is 100, and the statistics collection and comparison interval is 10 seconds for input error packet alarms.

### Views

System view

### Predefined user roles

network-admin

### Parameters

**high-threshold** *high-value*: Specifies the upper threshold for input error packet alarms, in the range of 1 to 4294967295 packets.

**low-threshold** *low-value*: Specifies the lower threshold for input error packet alarms, in the range of 1 to 4294967295 packets.

**interval** *interval*: Specifies the statistics collection and comparison interval for input error packets, in the range of 1 to 65535 seconds.

**shutdown**: Shuts down an interface when the number of input error packets on the interface exceeds the upper threshold. Then, the interface stops forwarding all packets. To recover the interface, execute the **undo shutdown** command on the interface. If you do not specify this keyword, an upper threshold exceeding alarm is generated and the interface enters the alarm state when the number of input error packets exceeds the upper threshold on the interface.

**slot** *slot-number*: Specifies an IRF member device by its member ID.

### Usage guidelines

With the input error packet alarm function enabled, when the number of input error packets on an interface in normal state within the specified interval exceeds the upper threshold, the interface generates an upper threshold exceeding alarm and enters the alarm state. When the number of input error packets on an interface in the alarm state within the specified interval drops below the lower threshold, the interface generates a recovery alarm and restores to the normal state.

You can configure the input error packet alarm parameters in system view and interface view.

- The configuration in system view takes effect on all interfaces of the specified slot. The configuration in interface view takes effect only on the current interface.
- For an interface, the configuration in interface view takes priority, and the configuration in system view is used only when no configuration is made in interface view.

When you execute this command multiple times, the most recent configuration takes effect.

### Examples

# Set the upper threshold to 5000, lower threshold to 400, and statistics collection and comparison interval to 6 seconds for input error packet alarms.

```
<Sysname> system-view
```

```
[Sysname] ifmonitor input-error slot 1 high-threshold 5000 low-threshold 400 interval 6
```

## Related commands

**snmp-agent trap enable ifmonitor**

## ifmonitor input-usage

Use **ifmonitor input-usage** to configure global input bandwidth usage alarm parameters.

Use **undo ifmonitor input-usage** to restore the default.

## Syntax

```
ifmonitor input-usage slot slot-number high-threshold high-value
low-threshold low-value

undo ifmonitor input-usage slot slot-number
```

## Default

The upper threshold is 90, and the lower threshold is 80.

## Views

System view

## Predefined user roles

network-admin

## Parameters

**high-threshold** *high-value*: Specifies the upper threshold for input bandwidth usage alarms, in the range of 1 to 100.

**low-threshold** *low-value*: Specifies the lower threshold for input bandwidth usage alarms, in the range of 1 to 100.

**slot** *slot-number*: Specifies an IRF member device by its member ID.

## Usage guidelines

With the input bandwidth usage alarm function enabled, when the input bandwidth usage on an interface in normal state within the most recent statistics polling interval exceeds the upper threshold, the interface generates an upper threshold exceeding alarm and enters the alarm state. When the input bandwidth usage on an interface in the alarm state within the most recent statistics polling interval drops below the lower threshold, the interface generates a recovery alarm and restores to the normal state.

To set the statistics polling interval, use the **flow-interval** command.

You can configure the input bandwidth usage alarm parameters in system view and interface view.

- The configuration in system view takes effect on all interfaces of the specified slot. The configuration in interface view takes effect only on the current interface.
- For an interface, the configuration in interface view takes priority, and the configuration in system view is used only when no configuration is made in interface view.

For this command to take effect, you must use the **snmp-agent trap enable ifmonitor** command to enable the input bandwidth usage alarm function.

When you execute this command multiple times, the most recent configuration takes effect.

## Examples

# Set the upper threshold to 95 and lower threshold to 80 for input bandwidth usage alarms.

```
<Sysname> system-view
```

```
[Sysname] ifmonitor input-usage slot 1 high-threshold 95 low-threshold 80
```

## Related commands

`flow-interval`  
`snmp-agent trap enable ifmonitor`

## ifmonitor output-error

Use `ifmonitor output-error` to configure global output error packet alarm parameters.

Use `undo ifmonitor output-error` to restore the default.

## Syntax

```
ifmonitor output-error slot slot-number high-threshold high-value
low-threshold low-value interval interval [shutdown]
undo ifmonitor output-error slot slot-number
```

## Default

The upper threshold is 1000, the lower threshold is 100, and the statistics collection and comparison interval is 10 seconds for output error packet alarms.

## Views

System view

## Predefined user roles

network-admin

## Parameters

**high-threshold** *high-value*: Specifies the upper threshold for output error packet alarms, in the range of 1 to 4294967295 packets.

**low-threshold** *low-value*: Specifies the lower threshold for output error packet alarms, in the range of 1 to 4294967295 packets.

**interval** *interval*: Specifies the statistics collection and comparison interval for output error packets, in the range of 1 to 65535 seconds.

**shutdown**: Shuts down an interface when the number of output error packets on the interface exceeds the upper threshold. Then, the interface stops forwarding all packets. To recover the interface, execute the `undo shutdown` command on the interface. If you do not specify this keyword, an upper threshold exceeding alarm is generated and the interface enters the alarm state when the number of output error packets exceeds the upper threshold on the interface.

**slot** *slot-number*: Specifies an IRF member device by its member ID.

## Usage guidelines

With the output error packet alarm function enabled, when the number of output error packets on an interface in normal state within the specified interval exceeds the upper threshold, the interface generates an upper threshold exceeding alarm and enters the alarm state. When the number of output error packets on an interface in the alarm state within the specified interval drops below the lower threshold, the interface generates a recovery alarm and restores to the normal state.

You can configure the output error packet alarm parameters in system view and interface view.

- The configuration in system view takes effect on all interfaces of the specified slot. The configuration in interface view takes effect only on the current interface.
- For an interface, the configuration in interface view takes priority, and the configuration in system view is used only when no configuration is made in interface view.

When you execute this command multiple times, the most recent configuration takes effect.

## Examples

```
Set the upper threshold to 5000, lower threshold to 400, and statistics collection and comparison interval to 6 seconds for output error packet alarms.
```

```
<Sysname> system-view
```

```
[Sysname] ifmonitor output-error slot 1 high-threshold 5000 low-threshold 400 interval 6
```

## Related commands

```
snmp-agent trap enable ifmonitor
```

## ifmonitor output-usage

Use **ifmonitor output-usage** to configure global output bandwidth usage alarm parameters.

Use **undo ifmonitor output-usage** to restore the default.

## Syntax

```
ifmonitor output-usage slot slot-number high-threshold high-value
low-threshold low-value
```

```
undo ifmonitor output-usage slot slot-number
```

## Default

The upper threshold is 90, and the lower threshold is 80.

## Views

System view

## Predefined user roles

network-admin

## Parameters

**high-threshold** *high-value*: Specifies the upper threshold for output bandwidth usage alarms, in the range of 1 to 100.

**low-threshold** *low-value*: Specifies the lower threshold for output bandwidth usage alarms, in the range of 1 to 100.

**slot** *slot-number*: Specifies an IRF member device by its member ID.

## Usage guidelines

With the output bandwidth usage alarm function enabled, when the output bandwidth usage on an interface in normal state within the most recent statistics polling interval exceeds the upper threshold, the interface generates an upper threshold exceeding alarm and enters the alarm state. When the output bandwidth usage on an interface in the alarm state within the most recent statistics polling interval drops below the lower threshold, the interface generates a recovery alarm and restores to the normal state.

To set the statistics polling interval, use the **flow-interval** command.

You can configure the output bandwidth usage alarm parameters in system view and interface view.

- The configuration in system view takes effect on all interfaces of the specified slot. The configuration in interface view takes effect only on the current interface.
- For an interface, the configuration in interface view takes priority, and the configuration in system view is used only when no configuration is made in interface view.

For this command to take effect, you must use the **snmp-agent trap enable ifmonitor** command to enable the output bandwidth usage alarm function.

When you execute this command multiple times, the most recent configuration takes effect.

## Examples

# Set the upper threshold to 80 and lower threshold to 60 for output bandwidth usage alarms.

```
<Sysname> system-view
```

```
[Sysname] ifmonitor output-usage slot 1 high-threshold 80 low-threshold 60
```

## Related commands

**flow-interval**

**snmp-agent trap enable ifmonitor**

## port ifmonitor crc-error

Use **port ifmonitor crc-error** to configure CRC error packet alarm parameters for an interface.

Use **undo port ifmonitor crc-error** to restore the default.

## Syntax

```
port ifmonitor crc-error [ratio] high-threshold high-value low-threshold low-value interval interval [shutdown]
```

```
undo port ifmonitor crc-error
```

## Default

An interface uses the global CRC error packet alarm parameters.

## Views

Ethernet interface view

## Predefined user roles

network-admin

## Parameters

**ratio**: Specifies the alarm thresholds in percentage. If you do not specify this keyword, you configure the alarm thresholds in absolute value.

**high-threshold** *high-value*: Specifies the upper threshold for CRC error packet alarms. If you specify the **ratio** keyword, the value range is 1 to 100. If you do not specify the **ratio** keyword, the value range is 1 to 4294967295 packets.

**low-threshold** *low-value*: Specifies the lower threshold for CRC error packet alarms. If you specify the **ratio** keyword, the value range is 1 to 100. If you do not specify the **ratio** keyword, the value range is 1 to 4294967295 packets.

**interval** *interval*: Specifies the statistics collection and comparison interval for CRC error packets, in the range of 1 to 65535 seconds.

**shutdown**: Shuts down an interface when the number of incoming CRC error packets on the interface exceeds the upper threshold. Then, the interface stops forwarding all packets. To recover the interface, execute the **undo shutdown** command on the interface. If you do not specify this keyword, an upper threshold exceeding alarm is generated and the interface enters the alarm state when the number of incoming CRC error packets exceeds the upper threshold on the interface.

## Usage guidelines

With the CRC error packet alarm function enabled, when the number of incoming CRC error packets on an interface in normal state within the specified interval exceeds the upper threshold, the interface generates an upper threshold exceeding alarm and enters the alarm state. When the number of

incoming CRC error packets on an interface in the alarm state within the specified interval drops below the lower threshold, the interface generates a recovery alarm and restores to the normal state.

You can configure the CRC error packet alarm parameters in system view and interface view.

- The configuration in system view takes effect on all interfaces of the specified slot. The configuration in interface view takes effect only on the current interface.
- For an interface, the configuration in interface view takes priority, and the configuration in system view is used only when no configuration is made in interface view.

When you execute this command multiple times, the most recent configuration takes effect.

## Examples

# Set the upper threshold to 5000, lower threshold to 400, and statistics collection and comparison interval to 6 seconds for CRC error packet alarms on Ten-GigabitEthernet 1/0/1.

```
<Sysname> system-view
[Sysname] interface Ten-GigabitEthernet 1/0/1
[Sysname-Ten-GigabitEthernet1/0/1] port ifmonitor crc-error high-threshold 5000
low-threshold 400 interval 6
```

## Related commands

**snmp-agent trap enable ifmonitor**

## port ifmonitor input-error

Use **port ifmonitor input-error** to configure input error packet alarm parameters for an interface.

Use **undo port ifmonitor input-error** to restore the default.

## Syntax

**port ifmonitor input-error high-threshold** *high-value* **low-threshold** *low-value* **interval** *interval* [**shutdown**]

**undo port ifmonitor input-error**

## Default

An interface uses the global input error packet alarm parameters.

## Views

Ethernet interface view

## Predefined user roles

network-admin

## Parameters

**high-threshold** *high-value*: Specifies the upper threshold for input error packet alarms, in the range of 1 to 4294967295 packets.

**low-threshold** *low-value*: Specifies the lower threshold for input error packet alarms, in the range of 1 to 4294967295 packets.

**interval** *interval*: Specifies the statistics collection and comparison interval for input error packets, in the range of 1 to 65535 seconds.

**shutdown**: Shuts down an interface when the number of input error packets on the interface exceeds the upper threshold. Then, the interface stops forwarding all packets. To recover the interface, execute the **undo shutdown** command on the interface. If you do not specify this keyword, an upper threshold exceeding alarm is generated and the interface enters the alarm state when the number of input error packets exceeds the upper threshold on the interface.

## Usage guidelines

With the input error packet alarm function enabled, when the number of input error packets on an interface in normal state within the specified interval exceeds the upper threshold, the interface generates an upper threshold exceeding alarm and enters the alarm state. When the number of input error packets on an interface in the alarm state within the specified interval drops below the lower threshold, the interface generates a recovery alarm and restores to the normal state.

You can configure the input error packet alarm parameters in system view and interface view.

- The configuration in system view takes effect on all interfaces of the specified slot. The configuration in interface view takes effect only on the current interface.
- For an interface, the configuration in interface view takes priority, and the configuration in system view is used only when no configuration is made in interface view.

When you execute this command multiple times, the most recent configuration takes effect.

## Examples

# Set the upper threshold to 5000, lower threshold to 400, and statistics collection and comparison interval to 6 seconds for input error packet alarms on Ten-GigabitEthernet 1/0/1.

```
<Sysname> system-view
[Sysname] interface Ten-GigabitEthernet 1/0/1
[Sysname-Ten-GigabitEthernet1/0/1] port ifmonitor input-error high-threshold 5000
low-threshold 400 interval 6
```

## Related commands

**snmp-agent trap enable ifmonitor**

## port ifmonitor input-usage

Use **port ifmonitor input-usage** to configure input bandwidth usage alarm parameters.

Use **undo port ifmonitor input-usage** to restore the default.

## Syntax

```
port ifmonitor input-usage high-threshold high-value low-threshold
low-value
undo port ifmonitor input-usage
```

## Default

An interface uses the global input bandwidth usage alarm parameters.

## Views

Interface view

## Predefined user roles

network-admin

## Parameters

**high-threshold** *high-value*: Specifies the upper threshold for input bandwidth usage alarms, in the range of 1 to 100.

**low-threshold** *low-value*: Specifies the lower threshold for input bandwidth usage alarms, in the range of 1 to 100.

## Usage guidelines

With the input bandwidth usage alarm function enabled, when the input bandwidth usage on an interface in normal state within the most recent statistics polling interval exceeds the upper threshold, the interface generates an upper threshold exceeding alarm and enters the alarm state. When the

input bandwidth usage on an interface in the alarm state within the most recent statistics polling interval drops below the lower threshold, the interface generates a recovery alarm and restores to the normal state.

To set the statistics polling interval, use the **flow-interval** command.

You can configure the input bandwidth usage alarm parameters in system view and interface view.

- The configuration in system view takes effect on all interfaces of the specified slot. The configuration in interface view takes effect only on the current interface.
- For an interface, the configuration in interface view takes priority, and the configuration in system view is used only when no configuration is made in interface view.

For this command to take effect, you must use the **snmp-agent trap enable ifmonitor** command to enable the input bandwidth usage alarm function.

When you execute this command multiple times, the most recent configuration takes effect.

## Examples

```
Set the upper threshold to 80 and lower threshold to 60 for input bandwidth usage alarms on Ten-GigabitEthernet 1/0/1.
```

```
<Sysname> system-view
```

```
[Sysname] interface ten-gigabitethernet 1/0/1
```

```
[Sysname-Ten-GigabitEthernet1/0/1] port ifmonitor input-usage high-threshold 80
low-threshold 60
```

## Related commands

**flow-interval**

**snmp-agent trap enable ifmonitor**

## port ifmonitor output-error

Use **port ifmonitor output-error** to configure output error packet alarm parameters for an interface.

Use **undo port ifmonitor output-error** to restore the default.

## Syntax

```
port ifmonitor output-error high-threshold high-value low-threshold low-value interval interval [shutdown]
```

```
undo port ifmonitor output-error
```

## Default

An interface uses the global output error packet alarm parameters.

## Views

Ethernet interface view

## Predefined user roles

network-admin

## Parameters

**high-threshold** *high-value*: Specifies the upper threshold for output error packet alarms, in the range of 1 to 4294967295 packets.

**low-threshold** *low-value*: Specifies the lower threshold for output error packet alarms, in the range of 1 to 4294967295 packets.

**interval** *interval*: Specifies the statistics collection and comparison interval for output error packets, in the range of 1 to 65535 seconds.

**shutdown**: Shuts down an interface when the number of output error packets on the interface exceeds the upper threshold. Then, the interface stops forwarding all packets. To recover the interface, execute the **undo shutdown** command on the interface. If you do not specify this keyword, an upper threshold exceeding alarm is generated and the interface enters the alarm state when the number of output error packets exceeds the upper threshold on the interface.

## Usage guidelines

With the output error packet alarm function enabled, when the number of output error packets on an interface in normal state within the specified interval exceeds the upper threshold, the interface generates an upper threshold exceeding alarm and enters the alarm state. When the number of output error packets on an interface in the alarm state within the specified interval drops below the lower threshold, the interface generates a recovery alarm and restores to the normal state.

You can configure the output error packet alarm parameters in system view and interface view.

- The configuration in system view takes effect on all interfaces of the specified slot. The configuration in interface view takes effect only on the current interface.
- For an interface, the configuration in interface view takes priority, and the configuration in system view is used only when no configuration is made in interface view.

When you execute this command multiple times, the most recent configuration takes effect.

## Examples

# Set the upper threshold to 5000, lower threshold to 400, and statistics collection and comparison interval to 6 seconds for output error packet alarms on Ten-GigabitEthernet 1/0/1.

```
<Sysname> system-view
[Sysname] interface Ten-GigabitEthernet 1/0/1
[Sysname-Ten-GigabitEthernet1/0/1] port ifmonitor output-error high-threshold 5000
low-threshold 400 interval 6
```

## Related commands

**snmp-agent trap enable ifmonitor**

## port ifmonitor output-usage

Use **port ifmonitor output-usage** to configure output bandwidth usage alarm parameters.

Use **undo port ifmonitor output-usage** to restore the default.

## Syntax

```
port ifmonitor output-usage high-threshold high-value low-threshold low-value
undo port ifmonitor output-usage
```

## Default

An interface uses the global output bandwidth usage alarm parameters.

## Views

Interface view

## Predefined user roles

network-admin

## Parameters

**high-threshold** *high-value*: Specifies the upper threshold for output bandwidth usage alarms, in the range of 1 to 100.

**low-threshold** *low-value*: Specifies the lower threshold for output bandwidth usage alarms, in the range of 1 to 100.

## Usage guidelines

With the output bandwidth usage alarm function enabled, when the output bandwidth usage on an interface in normal state within the most recent statistics polling interval exceeds the upper threshold, the interface generates an upper threshold exceeding alarm and enters the alarm state. When the output bandwidth usage on an interface in the alarm state within the most recent statistics polling interval drops below the lower threshold, the interface generates a recovery alarm and restores to the normal state.

To set the statistics polling interval, use the **flow-interval** command.

You can configure the output bandwidth usage alarm parameters in system view and interface view.

- The configuration in system view takes effect on all interfaces of the specified slot. The configuration in interface view takes effect only on the current interface.
- For an interface, the configuration in interface view takes priority, and the configuration in system view is used only when no configuration is made in interface view.

For this command to take effect, you must use the **snmp-agent trap enable ifmonitor** command to enable the output bandwidth usage alarm function.

When you execute this command multiple times, the most recent configuration takes effect.

## Examples

# Set the upper threshold to 9 and lower threshold to 7 for output bandwidth usage alarms on Ten-GigabitEthernet 1/0/1.

```
<Sysname> system-view
[Sysname] interface ten-gigabitethernet 1/0/1
[Sysname-Ten-GigabitEthernet1/0/1] port ifmonitor output-usage high-threshold 9
low-threshold 7
```

## Related commands

**flow-interval**

**snmp-agent trap enable ifmonitor**

## snmp-agent trap enable ifmonitor

Use **snmp-agent trap enable ifmonitor** to enable interface alarm functions.

Use **undo snmp-agent trap enable ifmonitor** to disable interface alarm functions.

## Syntax

```
snmp-agent trap enable ifmonitor [crc-error | input-error | input-usage
| output-error | output-usage] *
undo snmp-agent trap enable ifmonitor [crc-error | input-error |
input-usage | output-error | output-usage] *
```

## Default

Interface alarm functions are enabled.

## Views

System view

## Predefined user roles

network-admin

## Parameters

**crc-error**: Enables the CRC error packet alarm function for interfaces.

**input-error**: Enables the input error packet alarm function for interfaces.

**input-usage**: Enables the input bandwidth usage alarm function for interfaces.

**output-error**: Enables the output error packet alarm function for interfaces.

**output-usage**: Enables the output bandwidth usage alarm function for interfaces.

## Examples

# Enable the CRC error packet alarm function for interfaces.

```
<Sysname> system-view
```

```
[Sysname] snmp-agent trap enable ifmonitor crc-error
```

# New feature: Establishing neighbors through the secondary IP address of an interface

## Enabling OSPF to establish neighbors through the secondary IP address of an interface

### About this task

By default, OSPF uses the primary IP address of an interface to establish neighbors. You can configure this feature to enable OSPF to establish neighbors through the secondary IP address of an interface.

### Restrictions and guidelines

- If an interface has both primary and secondary addresses and you have advertised the network that contains the primary address in an area of an OSPF process, OSPF uses the primary address to establish neighbors. For OSPF to establish neighbors through the secondary IP addresses, you must advertise the networks to which these addresses belong in the same area of the same OSPF process.
- If an interface does not have a primary address but has multiple secondary addresses, OSPF uses the lowest secondary address to establish neighbors. For OSPF to establish neighbors through the other secondary addresses, you must advertise the networks to which these addresses belong in the same area of the same OSPF process.
- OSPF cannot use secondary addresses for neighbor establishment on a P2P link if the local and remote addresses of the link belong to different networks.

### Procedure

1. Enter system view.  
**system-view**
2. Enter OSPF view.  
**ospf** [ *process-id* | **router-id** *router-id* | **vpn-instance** *vpn-instance-name* ] \*
3. Enter interface view.  
**interface** *interface-type* *interface-number*

4. Enable OSPF to establish neighbors through the secondary IP address of the interface.  
`ospf peer sub-address enable`

## Command reference

### New command: ospf peer sub-address enable

Use `ospf peer sub-address enable` to enable OSPF to establish neighbors through the secondary IP address of an interface.

Use `undo ospf peer sub-address enable` to disable OSPF from establishing neighbors through the secondary IP address of an interface.

#### Syntax

```
ospf peer sub-address enable
undo ospf peer sub-address enable
```

#### Default

OSPF cannot establish neighbors through the secondary IP address of an interface.

#### Views

Interface

#### Predefined user roles

network-admin

#### Usage guidelines

By default, OSPF uses the primary IP address of an interface to establish neighbors. You can execute this command to enable OSPF to establish neighbors through both the primary and secondary IP addresses of an interface.

If an interface has both primary and secondary addresses and you have advertised the network that contains the primary address in an area of an OSPF process, OSPF uses the primary address to establish neighbors. For OSPF to establish neighbors through the secondary IP addresses, you must advertise the networks to which these addresses belong in the same area of the same OSPF process.

If an interface does not have a primary address but has multiple secondary addresses, OSPF uses the lowest secondary address to establish neighbors. For OSPF to establish neighbors through the other secondary addresses, you must advertise the networks to which these addresses belong in the same area of the same OSPF process.

OSPF cannot use secondary addresses for neighbor establishment on a P2P link if the local and remote addresses of the link belong to different networks.

#### Examples

# On interface VLAN-interface 10, enable OSPF to establish neighbors through the secondary IP address of the interface.

```
<Sysname> system-view
[Sysname] interface vlan-interface 10
[Sysname-Vlan-interface10] ospf peer sub-address enable
```

#### Related commands

```
network
ospf area
```

Modified command: display ospf interface

### Syntax

```
display ospf [process-id] interface [interface-type interface-number |
verbose]
```

### Change description

The **Interface Address Type** field is added to the output of the **display ospf interface** command to display whether a neighbor is established through the primary or secondary IP address.

## New feature: Advertising only the global unicast address in the NEXT\_HOP attribute

### Advertising only the global unicast address in the NEXT\_HOP attribute

#### About this task

An IPv6 peer might fail to learn routes if it cannot parse a route update that contains both the link-local address and the global unicast address. To resolve this issue, perform this task to enable the local device to advertise only the global unicast address in the NEXT\_HOP attribute to its IPv6 peers.

#### Restrictions and guidelines

This feature might not apply to EBGP peers established through directly connected broadcast interfaces. If the next hop of the advertised route and the directly connected broadcast interfaces belong to the same subnet, this feature does not take effect.

#### Procedure (IPv4 unicast address family)

1. Enter system view.  
**system-view**
2. Enter BGP IPv4 unicast address family view or BGP-VPN IPv4 unicast address family view.
  - o Execute the following commands in sequence to enter BGP IPv4 unicast address family view:  
**bgp as-number [ instance instance-name ]**  
**address-family ipv4 [ unicast ]**
  - o Execute the following commands in sequence to enter BGP-VPN IPv4 unicast address family view:  
**bgp as-number [ instance instance-name ]**  
**ip vpn-instance vpn-instance-name**  
**address-family ipv4 [ unicast ]**
3. Enable the device to advertise only the global unicast address in the NEXT\_HOP attribute to its IPv6 peers.  
**nexthop global-address-only**  
By default, the local device with a link-local address advertises both the link-local address and the global unicast address in the NEXT\_HOP attribute to IPv6 BGP peers.

#### Procedure (IPv6 unicast/multicast address family)

1. Enter system view.

### **system-view**

2. Enter BGP IPv6 unicast address family view, BGP-VPN IPv6 unicast address family view, or BGP IPv6 multicast address family view.
  - Execute the following commands in sequence to enter BGP IPv6 unicast address family view:

```
bgp as-number [instance instance-name]
address-family ipv6 [unicast]
```
  - Execute the following commands in sequence to enter BGP-VPN IPv6 unicast address family view:

```
bgp as-number [instance instance-name]
ip vpn-instance vpn-instance-name
address-family ipv6 [unicast]
```
  - Execute the following commands in sequence to enter BGP IPv6 multicast address family view:

```
bgp as-number [instance instance-name]
address-family ipv6 multicast
```
3. Enable the device to advertise only the global unicast address in the NEXT\_HOP attribute to its IPv6 peers.

```
nexthop global-address-only
```

By default, the local device with a link-local address advertises both the link-local address and the global unicast address in the NEXT\_HOP attribute to IPv6 BGP peers.

## Command reference

### nexthop global-address-only

Use **nexthop global-address-only** to enable the device to advertise only the global unicast address in the NEXT\_HOP attribute to its IPv6 peers.

Use **undo nexthop global-address-only** to restore the default.

#### Syntax

```
nexthop global-address-only
undo nexthop global-address-only
```

#### Default

The local device with a link-local address advertises both the link-local address and the global unicast address in the NEXT\_HOP attribute to IPv6 BGP peers.

#### Views

BGP IPv4 unicast address family view  
BGP-VPN IPv4 unicast address family view  
BGP IPv6 unicast address family view  
BGP-VPN IPv6 unicast address family view  
BGP IPv6 multicast address family view

#### Predefined user roles

network-admin

## Usage guidelines

An IPv6 peer might fail to learn routes if it cannot parse a route update that contains both the link-local address and the global unicast address. To resolve this issue, execute this command to enable the local device to advertise only the global unicast address in the NEXT\_HOP attribute to its IPv6 peers.

This command might not apply to EBGp peers established through directly connected broadcast interfaces. If the next hop of the advertised route and the directly connected broadcast interfaces belong to the same subnet, this command does not take effect.

## Examples

# In BGP IPv6 unicast address family view, enable the device to advertise only the global unicast address in the NEXT\_HOP attribute to its IPv6 peers.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] address-family ipv6 unicast
[Sysname-bgp-default-ipv6] nexthop global-address-only
```

# New feature: Assigning a VXLAN tunnel interface to a VXLAN tunnel group

## Assigning a VXLAN tunnel interface to a VXLAN tunnel group

### About this task

This feature assigns a VXLAN tunnel interface to a VXLAN tunnel group. A VXLAN tunnel group contains one or multiple VXLAN tunnel interfaces on the same device.

Use this feature in conjunction with the traffic redirection feature to load share traffic among multiple VXLAN tunnels in a VXLAN tunnel group. For more information about traffic redirection, see "[Modified feature: Configuring a traffic redirecting action in a QoS traffic behavior.](#)"

### Restrictions and guidelines

A VXLAN tunnel interface can be assigned only to one VXLAN tunnel group. To assign the VXLAN tunnel interface to another VXLAN tunnel group, first remove the VXLAN tunnel interface from the original group by using the **undo group** command.

A VXLAN tunnel group can contain a maximum of 32 VXLAN tunnel interfaces.

### Procedure

1. Enter system view.  
**system-view**
  2. Enter VXLAN tunnel interface view.  
**interface tunnel** *number*
  3. Assign the VXLAN tunnel interface to a VXLAN tunnel group.  
**group** *group-id*
- By default, a VXLAN tunnel interface is not assigned to any VXLAN tunnel group.

## Command reference

### group

Use **group** to assign a VXLAN tunnel interface to a VXLAN tunnel group.

Use **undo group** to remove a VXLAN tunnel interface from its VXLAN tunnel group.

#### Syntax

```
group group-id
undo group
```

#### Default

A VXLAN tunnel interface is not assigned to any VXLAN tunnel group.

#### Views

VXLAN tunnel interface view

#### Predefined user roles

network-admin

#### Parameters

*group-id*: Specifies the ID of the VXLAN tunnel group, in the range of 1 to 32.

#### Usage guidelines

A VXLAN tunnel group can contain a maximum 32 VXLAN tunnel interfaces on the device.

Use this command in conjunction with the traffic redirection feature to load share traffic among multiple VXLAN tunnels in a VXLAN tunnel group. For more information about traffic redirection, see "[Modified feature: Configuring a traffic redirecting action in a QoS traffic behavior.](#)"

A VXLAN tunnel interface can be assigned only to one VXLAN tunnel group. To assign the VXLAN tunnel interface to another VXLAN tunnel group, first remove the VXLAN tunnel interface from the original group by using the **undo group** command.

#### Examples

```
Assign VXLAN tunnel interface 1 to VXLAN tunnel group 2.
<Sysname> system-view
[Sysname] interface tunnel 1 mode vxlan
[Sysname-Tunnel1] group 2
```

## New feature: Setting VLAN tags for sent protocol packets on an interface

### Enabling the function of setting VLAN tags for sent protocol packets on an interface

#### About this task

With this feature enabled on an interface, the interface records all layers of VLAN tags of received LACP or BFD protocol packets. When the interface sends the same protocol packets, the interface adds the recorded VLAN tags to the sent protocol packets, regardless of whether the link type of the interface allows these packets to carry these VLAN tags. For more information about LLDP, see

LLDP configuration in *Layer 2—LAN Switching Configuration Guide*. For more information about BFD, see BFD configuration in *High Availability Configuration Guide*.

## Restrictions and guidelines

Executing the **port outbound-vlan-tag enable** command and then its **undo** form on an interface will cause BFD sessions on the interface to go down. Perform this operation with caution.

Executing this command on an aggregation group member port does not affect the other member ports in the same aggregation group or affect the Selected state of aggregation group member ports. To ensure configuration consistency within an aggregation group, manually execute this command on all aggregation group member ports.

## Procedure

1. Enter system view.  
**system-view**
2. Enter Layer 2 Ethernet interface view.  
**interface** *interface-type* *interface-number*
3. Enable the function of setting VLAN tags for sent protocol packets on the interface.  
**port outbound-vlan-tag enable**

## Command reference

### port outbound-vlan-tag enable

Use **port outbound-vlan-tag enable** to enable the function of setting VLAN tags for sent protocol packets on an interface.

Use the **undo port outbound-vlan-tag enable** command to disable the function of setting VLAN tags for sent protocol packets on an interface.

### Syntax

```
port outbound-vlan-tag enable
undo port outbound-vlan-tag enable
```

### Default

The function of setting VLAN tags for sent protocol packets is disabled on an interface.

### Views

Layer 2 Ethernet interface view

### Predefined user roles

network-admin

### Usage guidelines

With this feature enabled on an interface, the interface records all layers of VLAN tags of received LACP or BFD protocol packets. When the interface sends the same protocol packets, the interface adds the recorded VLAN tags to the sent protocol packets, regardless of whether the link type of the interface allows these packets to carry these VLAN tags. For more information about LLDP, see LLDP configuration in *Layer 2—LAN Switching Configuration Guide*. For more information about BFD, see BFD configuration in *High Availability Configuration Guide*.

Executing the **port outbound-vlan-tag enable** command and then its **undo** form on an interface will cause BFD sessions on the interface to go down. Perform this operation with caution.

Executing this command on an aggregation group member port does not affect the other member ports in the same aggregation group or affect the Selected state of aggregation group member ports.

To ensure configuration consistency within an aggregation group, manually execute this command on all aggregation group member ports.

## Examples

# Enable the function of setting VLAN tags for sent protocol packets on Ten-GigabitEthernet1/0/1.

```
<System> system-view
[System] interfaceten-gigabitethernet 1/0/1
[System-Ten-GigabitEthernet1/0/1] port outbound-vlan-tag enable
```

# New feature: Configuring the system ID

## Configuring the system ID

### About this task

You can use the system ID to indicate the position or functionality of the device or any other information.

### Procedure

1. Enter system view.  
**system-view**
2. Configure the system ID.  
**sysid** *system-id*

By default, the device does not have a system ID.

## Command reference

### sysid

Use **sysid** to set the system ID.

Use **undo sysid** to restore the default.

### Syntax

```
sysid system-id
undo sysid
```

### Default

The device does not have a system ID.

### Views

System view

### Predefined user roles

network-admin

### Parameters

*system-id*: Specifies the system ID for the device. You can use this argument to indicate the position or functionality of the device or any other information.

### Usage guidelines

Non-default vSystems do not support this command.

## Examples

# Set the system ID of the device to **position-hall**.

```
<Sysname> system-view
```

```
[Sysname] sysid positon-hall
```

# New feature: IPv6 duplicate detection on duplicate addresses

## Enabling duplicate detection on duplicate addresses

### About this task

If the system detects that an IPv6 address on an interface has been used on the network, the device marks that IPv6 address as duplicate. The interface cannot use the address for communication.

By default, an interface does not perform duplicate detection on duplicate addresses. Once an IPv6 address is marked as duplicate on an interface, it will be unusable even after it becomes unique on the link later.

To resolve this issue, enable duplicate detection for duplicate addresses. This feature regularly sends NS messages to the duplicate address until it does not receive an NA response message from that address or until duplicate detection is disabled for duplicate addresses.

You can set the maximum duplicate detection interval for duplicate addresses. After the device marks a detected address as duplicate, it waits for a random amount of time between 1 and the maximum detection interval. Then, the device resends an NS message to the solicited-node multicast address of the duplicate address. This mechanism helps reduce the risk of congestion that results from the NS messages sent for duplicate detection.

### Procedure

1. Enter system view.  
**system-view**
2. Enable duplicate detection on duplicate addresses.  
**ipv6 address duplicate-detect enable**  
By default, duplicate detection is disabled on duplicate addresses.
3. (Optional.) Set the maximum duplicate detection interval for duplicate addresses.  
**ipv6 address duplicate-detect interval *interval***  
By default, the maximum duplicate detection interval for duplicate addresses is 5 seconds.

## Command reference

### ipv6 address duplicate-detect enable

Use **ipv6 address duplicate-detect enable** to enable duplicate detection on duplicate addresses.

Use **undo ipv6 address duplicate-detect enable** to disable duplicate detection on duplicate addresses.

### Syntax

```
ipv6 address duplicate-detect enable
```

```
undo ipv6 address duplicate-detect enable
```

## Default

Duplicate detection is disabled on duplicate addresses.

## Views

System view

## Predefined user roles

network-admin

## Usage guidelines

If the system detects that an IPv6 address on an interface has been used on the network, the device marks that IPv6 address as duplicate. The interface cannot use the address for communication.

By default, an interface does not perform duplicate detection for duplicate addresses. Once an IPv6 address is marked as duplicate on an interface, it will be unusable even after it becomes unique on the link later.

To resolve this issue, enable duplicate detection on duplicate addresses. This feature regularly sends NS messages to the duplicate address until it does not receive an NA response message from that address or until duplicate detection is disabled on duplicate addresses.

To set the maximum duplicate detection interval for duplicate addresses, use the **ipv6 address duplicate-detect interval** command.

For more information about duplicate address detection, see IPv6 basics configuration in *Layer 3—IP Services Configuration Guide*.

## Examples

```
Enable duplicate detection on duplicate addresses.
<Sysname> system-view
[Sysname] ipv6 address duplicate-detect enable
```

## Related commands

**ipv6 address duplicate-detect interval**

## ipv6 address duplicate-detect interval

Use **ipv6 address duplicate-detect interval** to set the maximum duplicate detection interval for duplicate addresses.

Use **undo ipv6 address duplicate-detect interval** to restore the default.

## Syntax

```
ipv6 address duplicate-detect interval interval
undo ipv6 address duplicate-detect interval
```

## Default

The maximum duplicate detection interval for duplicate addresses is 5 seconds.

## Views

System view

## Predefined user roles

network-admin

## Parameters

*interval*: Sets the maximum duplicate detection interval for duplicate addresses in seconds. The value range for this argument is 1 to 60.

## Usage guidelines

You can set the maximum duplicate detection interval for duplicate addresses. After the device marks a detected address as duplicate, it waits for a random amount of time between 1 and the maximum detection interval. Then, the device resends an NS message to the solicited-node multicast address of the duplicate address. This mechanism helps reduce the risk of congestion that results from the NS messages sent for duplicate detection.

## Examples

```
Set the maximum duplicate detection interval to 10 seconds for duplicate addresses.
```

```
<Sysname> system-view
```

```
[Sysname] ipv6 address duplicate-detect interval 10
```

## Related commands

```
ipv6 address duplicate-detect enable
```

# New feature: BGP route re-origination

## Configuring BGP route re-origination

### About this task

If the route target of a received VPNv4/VPNv6/EVPN route matches the import route target of the local VPN instance, the local VPN instance re-originate the route and advertises the re-originated route to VPNv4/VPNv6/EVPN peers.

### Configuring route re-origination

1. Enter system view.  
**system-view**
2. Enter BGP instance view.  
**bgp as-number [ *instance instance-name* ]**
3. Enter BGP-VPN instance view.  
**ip vpn-instance *vpn-instance-name***
4. Enter BGP-VPN IPv4 unicast address family view or BGP-VPN IPv6 unicast address family view.
  - o Enter BGP-VPN IPv4 unicast address family view.  
**address-family ipv4 [ *unicast* ]**
  - o Enter BGP-VPN IPv6 unicast address family view.  
**address-family ipv6 [ *unicast* ]**
5. Configure the VPN instance to re-originate the optimal route and advertise the re-originated route to VPNv4/VPNv6/EVPN peers.  
**advertise route-reoriginate [ *route-policy route-policy-name* ]**  
**[ *replace-rt* ]**  
By default, the VPN instance does not advertise a received route to VPNv4/VPNv6/EVPN peers.

### Advertising the re-originated EVPN routes to IBGP peers

1. Enter system view.  
**system-view**
2. Enter BGP instance view.

- ```
bgp as-number [ instance instance-name ]
```
3. Enter BGP EVPN address family view.
address-family l2vpn evpn
 4. Configure the VPN instance to advertise the re-originated EVPN routes to IBGP peers.
peer { *group-name* | *ipv4-address* [*mask-length*] | *ipv6-address* [*prefix-length*] } **advertise vpn-reoriginate ibgp**
- By default, the VPN instance does not advertise the re-originated EVPN routes to IBGP peers.

Command reference

advertise route-reoriginate

Use **advertise route-reoriginate** to configure a VPN instance to re-originate the optimal route and advertise the re-originated route to VPNv4/VPNv6/EVPN peers.

Use **undo advertise route-reoriginate** to cancel the configuration.

Syntax

```
advertise route-reoriginate [ route-policy route-policy-name ]
[ replace-rt ]

undo advertise route-reoriginate
```

Default

A VPN instance does not advertise a received route to VPNv4/VPNv6/EVPN peers.

Views

BGP-VPN IPv4 unicast address family view

BGP-VPN IPv6 unicast address family view

Predefined user roles

network-admin

Parameters

route-policy *route-policy-name*: Specifies a routing policy to filter VPNv4/VPNv6/EVPN routes to be advertised to VPNv4/VPNv6/EVPN peers. The *route-policy-name* argument specifies a routing policy by its name, a case-sensitive string of 1 to 63 characters. If you do not specify this option, the command advertises all VPNv4/VPNv6/EVPN routes to VPNv4/VPNv6/EVPN peers.

replace-rt: Uses the route target of the local VPN instance to replace the route target of the VPNv4/VPNv6/EVPN route. If you do not specify this keyword, the route target of the route will not be changed.

Usage guidelines

If the route target of a received VPNv4/VPNv6/EVPN route matches the import route target of the local VPN instance, the local VPN instance re-originate the route and advertises the re-originated route to VPNv4/VPNv6/EVPN peers.

Examples

In BGP-VPN IPv4 unicast address family view, configure VPN instance **vpn1** to re-originate the optimal route and advertise the re-originated route to VPNv4 peers

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] ip vpn-instance vpn1
```

```
[Sysname-bgp-default-vpn1] address-family ipv4
[Sysname-bgp-default-ipv4-vpn1] advertise route-reoriginate
```

peer advertise vpn-reoriginate ibgp

Use **peer advertise vpn-reoriginate ibgp** to configure a VPN instance to advertise the re-originated EVPN routes to IBGP peers.

Use **undo peer advertise vpn-reoriginate ibgp** to restore the default.

Syntax

```
peer { group-name | ipv4-address [ mask-length ] | ipv6-address
[ prefix-length ] } advertise vpn-reoriginate ibgp
undo peer advertise vpn-reoriginate ibgp
```

Default

A VPN instance does not advertise the re-originated EVPN routes to IBGP peers.

Views

BGP EVPN address family view

Predefined user roles

network-admin

Parameters

group-name: Specifies a peer group by its name, a case-sensitive string of 1 to 47 characters. The specified peer group must have been created.

ipv4-address: Specifies a peer by its IPv4 address. The specified peer must have been created.

mask-length: Specifies a mask length in the range of 0 to 32. You can use the *ipv4-address* and *mask-length* arguments together to specify a subnet. If you specify a subnet, this command specifies all dynamic peers in the subnet.

ipv6-address: Specifies a peer by its IPv6 address. The specified peer must have been created.

prefix-length: Specifies a prefix length in the range of 0 to 128. You can use the *ipv6-address* and *prefix-length* arguments together to specify a subnet. If you specify a subnet, this command specifies all dynamic peers in the subnet.

Usage guidelines

After you execute the **advertise route-reoriginate** command, if the route target of a received EVPN route matches the import route target of the local VPN instance, the local VPN instance re-originates the route and advertises the re-originated route to EVPN peers. If an EVPN peer is established through IBGP, the local VPN instance does not advertise the re-originated route to the peer. To resolve this issue, execute this command.

Examples

In BGP EVPN address family view, configure the VPN instance to advertise the re-originated EVPN routes to IBGP peers.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] address-family l2vpn evpn
[Sysname-bgp-default-evpn] peer 1.1.1.1 advertise vpn-reoriginate ibgp
```

Related commands

advertise route-reoriginate

New feature: Sending RA messages over VXLAN tunnels

Enabling a distributed EVPN gateway to send RA messages over VXLAN tunnels

About this task

By default, a distributed EVPN gateway drops the RS messages received from VXLAN tunnels and periodically advertises RA messages only to the local site. As a result, a distributed EVPN gateway does not send RA messages over VXLAN tunnels, and remote gateways cannot update information about the gateway based on RA messages. To resolve the issue, perform this task to enable distributed EVPN gateways to reply to remote RS messages with RA messages and periodically advertise RA messages over VXLAN tunnels.

Procedure

1. Enter system view.
system-view
2. Globally enable VSI interfaces to send RA messages over VXLAN tunnels.
ipv6 nd ra tunnel-broadcast global enable
By default, VSI interfaces do not send RA messages over VXLAN tunnels.

Command reference

ipv6 nd ra tunnel-broadcast global enable

Use **ipv6 nd ra tunnel-broadcast global enable** to globally enable VSI interfaces to send RA messages over VXLAN tunnels.

Use **undo ipv6 nd ra tunnel-broadcast global enable** to globally disable VSI interfaces from sending RA messages over VXLAN tunnels.

Syntax

ipv6 nd ra tunnel-broadcast global enable
undo ipv6 nd ra tunnel-broadcast global enable

Default

VSI interfaces do not send RA messages over VXLAN tunnels.

Views

System view

Predefined user roles

network-admin

Usage guidelines

By default, a distributed EVPN gateway drops the RS messages received from VXLAN tunnels and periodically advertises RA messages only to the local site. As a result, a distributed EVPN gateway does not send RA messages over VXLAN tunnels, and remote gateways cannot update information about the gateway based on RA messages. To resolve the issue, use this command to enable distributed EVPN gateways to reply to remote RS messages with RA messages and periodically advertise RA messages over VXLAN tunnels.

Examples

```
# Globally enable VSI interfaces to send RA messages over VXLAN tunnels.  
<Sysname> system-view  
[Sysname] ipv6 nd ra tunnel-broadcast global enable
```

Related commands

```
ipv6 nd ra interval  
ipv6 nd ra tunnel-broadcast enable
```

New feature: SAVA

About SAVA

Source Address Validation Architecture (SAVA) checks the validity of the source IPv6 address of packets based on routing information to prevent source IPv6 address spoofing attacks.

A SAVA-enabled interface creates SAVA entries based on routes that use the interface as the outgoing interface. Upon receiving an IPv6 packet on the interface, the device searches for a SAVA entry that contains the source IPv6 address and the receiving interface. If a match is found, the device forwards the packet. If no match is found, the device discards the packet.

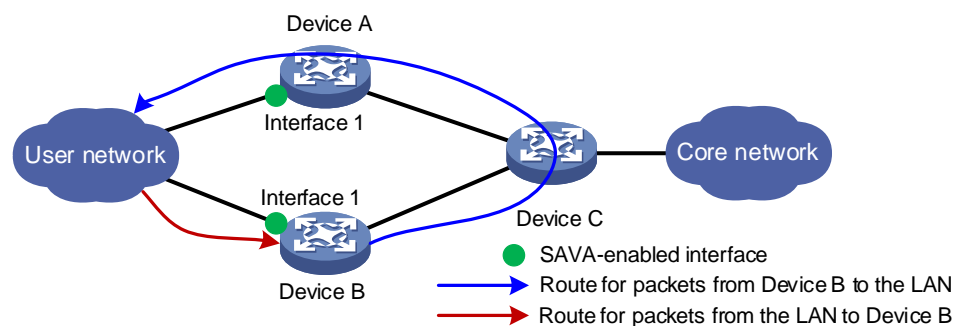
SAVA is typically deployed on the border devices of the backbone network connecting to an access network.

Benefits

As shown in [Figure 1](#), asymmetric routing might occur in an IPv6 multiple-access network. The outgoing interface for a return packet might be different from the incoming interface of the original packet. You can enable strict uRPF check on interface 1 of Device B to prevent attacks with spoofed source IPv6 addresses. However, strict uRPF check might determine that a valid packet from the LAN as an attack packet based on local routing information, which causes unexpected packet drop.

Compared with uRPF, SAVA can prevent source IPv6 address spoofing attacks and avoid mistaken packet drop caused by asymmetric routing. When SAVA is enabled on interface 1 of both Device A and Device B, the devices can create SAVA entries based on local routes and remote routes synchronized from other routing protocols. In this way, both the devices have consistent SAVA entries. This ensures that packets of valid users will not be mistakenly dropped on either a asymmetric network or a symmetric network with asymmetric routing.

Figure 1 SAVA-enabled network



Mechanism

The mechanism of SAVA includes obtaining prefixes of valid users, creating SAVA entries based on the prefixes, and using SAVA entries to filter IPv6 packets.

Obtaining prefixes of valid users

On an IPv6 multiple access network, each border device on the backbone network obtains prefix information of all users in the LAN based on the following types of routes:

- **Local routes**—Each border device obtains prefix information of users in the LAN from local routes, including direct routes, static routes, and dynamic routes for the interfaces connecting to the LAN.
- **Remote routes synchronized from other routing protocols**—Each border device adds the same tag to the local routes for the LAN and redistributes the routes to other border devices through a routing protocol. All the border devices import the remote routes with the specified tag and then obtain prefix information from these routes. In this way, all border devices have consistent prefix information of valid users.

Creating SAVA entries

Each border device creates SAVA entries for the SAVA-enabled interfaces based on prefix information of valid users. A SAVA entry includes a prefix, prefix length, and binding interface.

Filtering user packets based on SAVA entries

When a SAVA-enabled interface receives an IPv6 packet from the LAN, SAVA validates the source IPv6 address of the packet. It searches for a SAVA entry that includes the interface by the source IPv6 address. If a match is found, SAVA permits the packet. If no match is found, SAVA drops the packet.

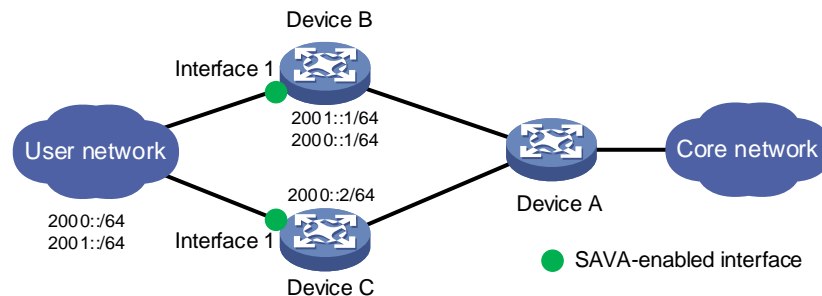
Application scenarios

Deployed on intra-AS border devices directly connected to the LAN

As shown in [Figure 2](#), users in the LAN connect to the backbone network through Device B and Device C. Device B is configured with two gateway addresses (2000::1/64 and 2001::1/64). Device C is configured with the gateway address (2000::2/64).

You can configure SAVA on interface 1 of both Device B and Device C. When Device C receives an IPv6 route with prefix 2001::/64 from Device B, it creates an SAVA entry with this prefix. Then, upon receiving an IPv6 packet from a user at 2001::/64, Device C will permit the packet because a matching SAVA entry exists.

Figure 2 SAVA deployed on border devices directly connected to the LAN

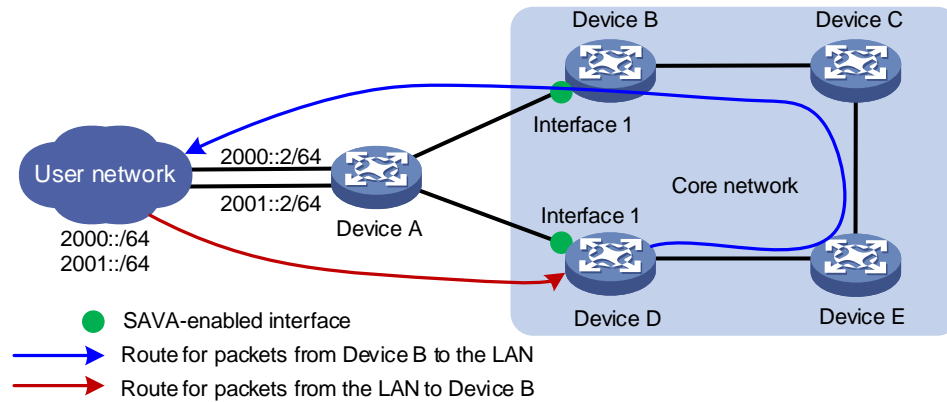


Deployed on intra-AS border devices indirectly connected to the LAN

As shown in [Figure 3](#), asymmetric traffic exists between Device D and the LAN after route convergence. Traffic from Device D to the LAN will traverse Device E, Device C, Device B, and Device A. Traffic from the LAN to Device D will traverse only Device A.

You can configure SAVA on interface 1 of both Device B and Device D. Device B and Device D synchronize routes with each other, obtain prefix information of the LAN based on the synchronized remote routes, and then create SAVA entries. Upon receiving a packet from a user in the LAN, Device C will permit the packet because a matching SAVA entry exists.

Figure 3 SAVA deployed on intra-AS border devices indirectly connected to the LAN

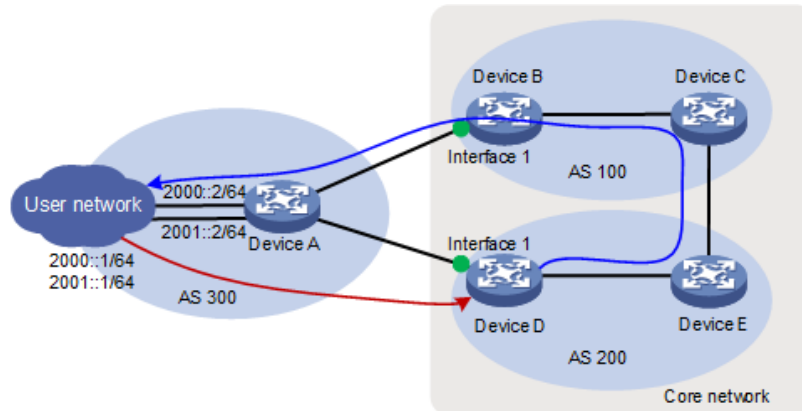


Deployed on inter-AS border devices indirectly connected to the LAN

As shown in Figure 4, asymmetric traffic exists between Device D and the LAN after route convergence. Traffic from Device D to the LAN will traverse Device E, Device C, Device B, and Device A. Traffic from the LAN to Device D will traverse only Device A.

You can configure SAVA on interface 1 of both Device B and Device D. Device B and Device C in AS 100 as well as Device D and Device E in AS 200 synchronize local routes with each other through an IGP protocol (OSPFv3, for example). Then, Device C in AS 100 and Device E in AS 200 synchronize routes with each other through BGP. Upon receiving a packet from a user in the LAN, Device C will permit the packet because a matching SAVA entry exists.

Figure 4 SAVA deployed on inter-AS border devices indirectly connected to the LAN



SAVA tasks at a glance

To configure SAVA, perform the following tasks:

1. [Enabling SAVA](#)
2. [Enabling SAVA entry creation based on synchronized remote routes](#)
 Perform this task if the LAN connects to the backbone network through multiple border devices and interfaces on the border devices do not have prefix information of all users in the LAN.
3. [Adding an interface to a SAVA access group](#)

Perform this task if a border device has multiple interfaces connected to the same LAN.

4. [Configuring SAVA logging](#)

Enabling SAVA

Restrictions and guidelines

If the device has a large number of routing entries, it might take a long time for the device to complete SAVA entry creation. Before SAVA entry creation completes, valid IPv6 packets might be dropped.

Procedure

1. Enter system view.
system-view
2. Enter interface view.
interface *interface-type interface-number*
3. Enable SAVA on the interface.
ipv6 sava enable
By default, SAVA is disabled.

Enabling SAVA entry creation based on synchronized remote routes

About this task

Perform this task to ensure that the border devices through which users in a LAN connects to the backbone network have the same SAVA entries to avoid mistaken packet drop.

Each border device adds a route tag to local routes based on which SAVA entries are created and then advertises the tagged local routes to the other border devices through a routing protocol. Then, other border devices will create SAVA entries upon receiving the tagged routes advertised by other border devices. In this way, the border devices maintain consistent SAVA entries.

Prerequisites

Before you enable SAVA entry creation based on synchronized remote routes, you must complete one of the following tasks:

- Configure OSPFv3 link tag inheritance
Perform this task if OSPFv3 runs on the network. For more information about OSPFv3 link tag inheritance, see OSPFv3 configuration in *Layer 3—IP Routing Configuration Guide*.
- Configure IPv6 IS-IS link tag inheritance
Perform this task if IPv6 IS-IS runs on the network. For more information about IPv6 IS-IS link tag inheritance, see IS-IS configuration in *Layer 3—IP Routing Configuration Guide*.

Procedure

1. Enter system view.
system-view
2. Enter interface view.
interface *interface-type interface-number*
3. Enable the device to create SAVA entries based on synchronized remote routes with the specified route tag.
ipv6 sava import remote-route-tag *tag*
By default, the device does not create SAVA entries based on synchronized remote routes.

Adding an interface to a SAVA access group

About this task

If the device has multiple interfaces connected to the same LAN, each interface creates SAVA entries only based on its local routes. When an interface receives a packet from the LAN for which the interface has no matching SAVA entry, the packet will be discarded.

To resolve this issue, you can add the interfaces to a SAVA access group. The interfaces in the same SAVA access group will synchronize SAVA entries created based on local routes with each other. This avoids unexpected packet drop caused by asymmetric routing.

Restrictions and guidelines

All interfaces in a SAVA access group must belong to the public network or the same VPN instance.

You can add a maximum of eight interfaces to a SAVA access group.

Procedure

1. Enter system view.
system-view
2. Enter interface view.
interface *interface-type* *interface-number*
3. Add the interface to a SAVA access group.
ipv6 sava access-group *group-name*
By default, an interface does not belong to a SAVA access group.

Configuring SAVA logging

About this task

To identify and troubleshoot issues, enable SAVA logging.

This feature enables the device to generate SAVA logs when SAVA detects spoofing packets.

With the information center, you can configure log destinations and output rules. For more information about the information center, see Network Management and Monitoring Configuration Guide.

Procedure

1. Enter system view.
system-view
2. Configure SAVA logging.
ipv6 sava log enable spoofing-packet [**interval** *interval* | **number** *number*] *****
By default, SAVA logging is disabled.

Display and maintenance commands for SAVA

Execute display commands in any view and reset commands in user view.

| Task | Command |
|-----------------------|--|
| Display SAVA entries. | display ipv6 sava [interface <i>interface-type</i> <i>interface-number</i>] [slot <i>slot-number</i>] |

| Task | Command |
|--------------------------------------|--|
| Display SAVA packet drop statistics. | display ipv6 sava packet-drop statistics [interface <i>interface-type interface-number</i>] |
| Clear SAVA packet drop statistics. | reset ipv6 sava packet-drop statistics [interface <i>interface-type interface-number</i>] |

Command reference

display ipv6 sava

Use **display ipv6 sava** to display SAVA entries.

Syntax

display ipv6 sava [**interface** *interface-type interface-number*] [**slot** *slot-number*]

Views

Any view

Predefined user roles

network-admin
network-operator

Parameters

interface *interface-type interface-number*: Specifies an interface by its type and number. If you do not specify an interface, this command displays SAVA entries for all interfaces.

slot *slot-number*: Specifies an IRF member device by its member ID. If you do not specify a member device, this command displays SAVA entries on the master device.

Examples

Display SAVA entries.

```
<Sysname> display ipv6 sava
```

```
IPv6 SAVA entry count: 2
```

```
Destination: 2011::
```

```
Prefix length: 64
```

```
Interface: Vlan-int10
```

```
Flags: L
```

```
Destination: 2012::
```

```
Prefix length: 64
```

```
Interface: Vlan-int20
```

```
Flags: L
```

Table 2 Command output

| Field | Description |
|-----------------------|---|
| IPv6 SAVA entry count | Number of SAVA entries. |
| Destination | Destination IPv6 address. |
| Prefix length | Prefix length of the IPv6 address. |
| Interface | Interface name. |
| Flag | Flag of the SAVA entry: <ul style="list-style-type: none"> L—Local entry. R—Remote entry. |

| Field | Description |
|-------|--|
| | <ul style="list-style-type: none"> G—Access group entry. |

display ipv6 sava packet-drop statistics

Use **display ipv6 sava packet-drop statistics** to display SAVA packet drop statistics.

Syntax

```
display ipv6 sava packet-drop statistics [ interface interface-type
interface-number ]
```

Views

Any view

Predefined user roles

network-admin

network-operator

Parameters

interface interface-type interface-number: Specifies an interface by its type and number. If you do not specify an interface, this command displays SAVA packet drop statistics for all interfaces.

Examples

Display SAVA packet drop statistics.

```
<Sysname> display ipv6 sava packet-drop statistics
Vlan-interface10:
  Packets:0                Bytes: 0

Vlan-interface20:
  Packets:10               Bytes: 1500
```

Table 3 Command output

| Field | Description |
|---------|------------------------------------|
| Packets | Number of packets dropped by SAVA. |
| Bytes | Number of bytes dropped by SAVA. |

Related commands

```
reset ipv6 sava packet-drop statistics
```

ipv6 sava access-group

Use **ipv6 sava access-group** to add an interface to an access group.

Use **undo ipv6 sava access-group** to remove an interface from an access group.

Syntax

```
ipv6 sava access-group group-name
undo ipv6 sava access-group
```

Default

An interface does not belong to any access group.

Views

Interface view

Predefined user roles

network-admin

Parameters

group-name: Specifies an access group by its name, a case-sensitive string of 1 to 255 characters.

Usage guidelines

If the device has multiple interfaces connected to the same LAN, the device might receive packets from users in the LAN on different interfaces. However, each interface creates SAVA entries only based on its local routes. When an interface receives a packet from the LAN for which the interface has no matching SAVA entry, the packet will be discarded.

To resolve this issue, you can add the interfaces to a SAVA access group. The interfaces in the SAVA access group will synchronize SAVA entries that are created based on local routes with each other. This avoids unexpected packet drop caused by asymmetric routing.

All interfaces in a SAVA access group must belong to the public network or the same VPN instance.

An interface can be added only to one SAVA access group. If you execute this command multiple times, the most recent configuration takes effect.

A SAVA access group can contain a maximum of eight interfaces.

Examples

Add VLAN-interface 10 to SAVA access group **aaa**.

```
<Sysname> system-view
```

```
[Sysname] interface vlan-interface 10
```

```
[Sysname-Vlan-interface10] ipv6 sava access-group aaa
```

Related commands

ipv6 sava enable

ipv6 sava enable

Use **ipv6 sava enable** to enable SAVA.

Use **undo ipv6 sava enable** to disable SAVA.

Syntax

ipv6 sava enable

undo ipv6 sava enable

Default

SAVA is disabled.

Views

Interface view

Predefined user roles

network-admin

Usage guidelines

If the device has a large number of routing entries, it might take a long time for the device to complete SAVA entry creation. Before SAVA entry creation completes, valid IPv6 packets might be dropped.

Examples

```
# Enable SAVA on VLAN-interface 10.
<Sysname> system-view
[Sysname] interface vlan-interface 10
[Sysname-Vlan-interface10] ipv6 sava enable
```

Related commands

```
display ipv6 sava
ipv6 sava access-group
```

ipv6 sava log enable spoofing-packet

Use **ipv6 sava log enable spoofing-packet** to enable SAVA logging.

Use **undo ipv6 sava log enable spoofing-packet** to disable SAVA logging.

Syntax

```
ipv6 sava log enable spoofing-packet [ interval interval | number number ]*
undo ipv6 sava log enable spoofing-packet
```

Default

SAVA logging is disabled.

Views

System view

Predefined user roles

network-admin

Parameters

interval *interval*: Specifies the interval at which the device outputs SAVA logs, in seconds. The value can be 0 or in the range of 5 to 3600, and the default is 60. If you set the interval to 0 seconds, the device outputs a SAVA log immediately after detecting an IPv6 source address spoofing packet.

number *number*: Specifies the maximum number of SAVA logs that can be outputted each time, in the range of 1 to 128. The default is 128.

Usage guidelines

To identify and troubleshoot issues, enable SAVA logging.

This feature enables the device to output SAVA logs when SAVA detects spoofing packets.

With the information center, you can configure log destinations and output rules. For more information about the information center, see Network Management and Monitoring Configuration Guide.

Outputting a large number of SAVA logs might degrade device performance and cause inconvenience for fault location. You can limit the number of SAVA logs that the device outputs each time.

An IRF member device can output a maximum of 128 SAVA logs each time.

Examples

```
# Enable SAVA logging.
<Sysname> system-view
[Sysname] ipv6 sava log enable spoofing-packet
```

ipv6 sava import remote-route-tag

Use **ipv6 sava import remote-route-tag** to enable an interface to create SAVA entries based on synchronized remote routes.

Use **undo ipv6 sava import remote-route-tag** to restore the default.

Syntax

```
ipv6 sava import remote-route-tag tag
undo ipv6 sava import remote-route-tag
```

Default

An interface does not create SAVA entries based on synchronized remote routes.

Views

Interface view

Predefined user roles

network-admin

Parameters

tag: Specifies a tag of synchronized remote routes, in the range of 1 to 4294967295.

Usage guidelines

This command enables an interface to create SAVA entries based on synchronized remote entries with the specified route tag.

Use this command if the LAN connects to the backbone network through multiple access devices and LAN-side interfaces on the border devices do not have prefix information of all users in the LAN. This task ensures that the border devices have the same SAVA entries to avoid mistaken packet drop.

Each border device adds a route tag to local routes based on which SAVA entries are created and then advertises the tagged local routes to the other border devices through a routing protocol. The other border devices will create SAVA entries upon receiving the tagged routes advertised by other border devices.

If you execute this command multiple times, the most recent configuration takes effect.

Examples

Configure the device to create SAVA entries based on synchronized remote entries with tag 10 on VLAN-interface 10.

```
<Sysname> system-view
[Sysname] interface vlan-interface 10
[Sysname-Vlan-interface10] ipv6 sava import remote-route-tag 100
```

reset ipv6 sava packet-drop statistics

Use **reset ipv6 sava packet-drop statistics** to clear SAVA packet drop statistics.

Syntax

```
reset ipv6 sava packet-drop statistics [ interface interface-type
interface-number ]
```

Views

Any view

Predefined user roles

network-admin

Parameters

interface *interface-type interface-number*: Specifies an interface by its type and number. If you do not specify an interface, this command clears SAVA packet drop statistics for all interfaces.

Examples

```
# Clear SAVA packet drop statistics.
<Sysname> reset ipv6 sava packet-drop statistics
```

Related commands

```
display ipv6 sava packet-drop statistics
```

New feature: IPv6 IS-IS link tag

Configuring the IPv6 IS-IS link tag feature

About this task

Set an IPv6 IS-IS link tag for a device (for example, Device A) and enable IPv6 IS-IS link tag inheritance on another device (for example, Device B) to filter routes as follows:

1. Device A advertises the IPv6 IS-IS link tag in LSPs.
2. After receiving the LSPs advertised by Device A, Device B inherits the IPv6 IS-IS link tag advertised by Device A.
If Device B receives multiple link tags from Device A, Device B selects and inherits only one tag.
3. Device B configures the link tag as the route tag for all routes that pass the link.
Then, the route tag can be used for route filtering.

This feature is usually used in Source Address Validation Architecture (SAVA) scenarios to ensure the consistency of IPv6 SAVA entries on multiple gateway devices:

1. Set an IPv6 IS-IS link tag for the customer-side interface on the gateway device that advertises IPv6 SAVA entries.
2. Enable IPv6 IS-IS link tag inheritance on the gateway device that receives IPv6 SAVA entries. When the route tag configured for a route is the same as that calculated based on the link tag, the device generates IPv6 SAVA entries based on the route prefix.

Setting an IPv6 IS-IS link tag for an interface

1. Enter system view.
system-view
2. Enter interface view.
interface *interface-type interface-number*
3. Set an IPv6 IS-IS link tag for the interface.
isis ipv6 link-tag *tag*

By default, no IPv6 IS-IS link tag is configured for the interface.

Enabling IPv6 IS-IS link tag inheritance

1. Enter system view.
system-view
2. Enter IS-IS view.
isis [*process-id*] [**vpn-instance** *vpn-instance-name*]
3. Enter IPv6 address family view.
address-family ipv6 [**unicast**]
4. Enable IPv6 IS-IS link tag inheritance.
link-tag inherit enable

By default, IPv6 IS-IS link tag inheritance is disabled.

Command reference

link-tag inherit enable

Use **link-tag inherit enable** to enable IPv6 IS-IS link tag inheritance.

Use **undo link-tag inherit enable** to disable IPv6 IS-IS link tag inheritance.

Syntax

```
link-tag inherit enable
undo link-tag inherit enable
```

Default

IPv6 IS-IS link tag inheritance is disabled.

Views

IS-IS IPv6 unicast address family view

Predefined user roles

network-admin

Usage guidelines

After you execute this command on a device, the device inherits the IPv6 IS-IS link tag in received LSPs. If the device receives multiple link tags from a device, it selects and inherits only one tag. Then, the device configures the link tag as the route tag for all routes that pass the link.

The IPv6 IS-IS link tag feature is usually used in Source Address Validation Architecture (SAVA) scenarios to ensure the consistency of IPv6 SAVA entries on multiple gateway devices:

1. Set an IPv6 IS-IS link tag for the customer-side interface on the gateway device that advertises IPv6 SAVA entries.
2. Enable IPv6 IS-IS link tag inheritance on the gateway device that receives IPv6 SAVA entries. The device generates an IPv6 SAVA entry based on a route prefix if the following two tags are the same:
 - The route tag configured by executing the **ipv6 sava import remote-route-tag** command.
 - The IPv6 IS-IS link tag that the route inherits.

Examples

In IS-IS process 1, enable IPv6 IS-IS link tag inheritance.

```
<Sysname> system-view
```

```
[Sysname] isis 1
[Sysname-isis-1] address-family ipv6
[Sysname-isis-1-ipv6] link-tag inherit enable
```

Related commands

isis ipv6 link-tag

isis ipv6 link-tag

Use **isis ipv6 link-tag** to set an IPv6 IS-IS link tag for an interface.

Use **undo isis ipv6 link-tag** to restore the default.

Syntax

```
isis ipv6 link-tag tag
undo isis ipv6 link-tag
```

Default

No IPv6 IS-IS link tag is configured for an interface.

Views

Interface view

Predefined user roles

network-admin

Parameters

tag: Specifies an IPv6 IS-IS link tag in the range of 1 to 4294967295.

Usage guidelines

Execute this command together with the **link-tag inherit enable** command to filter routes based on IPv6 IS-IS link tags.

If you execute this command multiple times on an interface, the most recent configuration takes effect.

Examples

```
# Set the IPv6 IS-IS link tag for interface VLAN-interface 10 to 222333.
<Sysname> system-view
[Sysname] interface vlan-interface 10
[Sysname-Vlan-interface10] isis ipv6 link-tag 222333
```

Related commands

link-tag inherit enable

New feature: OSPFv3 link tag

Configuring the OSPFv3 link tag feature

About this task

The interface link tag feature can work together with the link tag inheritance feature to filter routes. For example, you can set an OSPFv3 link tag for Device A and enable OSPFv3 link tag inheritance on Device B to filter routes as follows:

1. Device A advertises the OSPFv3 link tag in E-router LSAs.
2. After receiving the LSAs advertised by Device A, Device B inherits the OSPFv3 link tag advertised by Device A.
If Device B receives multiple link tags from Device A, Device B selects and inherits only one tag.
3. Device B configures the link tag as the route tag for all routes that pass the link.
Then, the route tag can be used for route filtering.

This feature is usually used in Source Address Validation Architecture (SAVA) scenarios to ensure the consistency of IPv6 SAVA entries on multiple gateway devices as follows:

1. Set an OSPFv3 link tag for the customer-side interface on the gateway device that advertises IPv6 SAVA entries.
2. Enable OSPFv3 link tag inheritance on the gateway device that receives IPv6 SAVA entries. The device generates an IPv6 SAVA entry based on a route prefix if the following two tags are the same:
 - The route tag configured by executing the **ipv6 sava import remote-route-tag** command.
 - The OSPFv3 link tag that the route inherits.

Setting an OSPFv3 link tag for an interface

1. Enter system view.
system-view
2. Enter interface view.
interface *interface-type* *interface-number*
3. Set an OSPFv3 link tag for the interface.
ospfv3 link-tag *tag* [**instance** *instance-id*]
By default, no OSPFv3 link tag is configured for the interface.

Enabling OSPFv3 link tag inheritance

1. Enter system view.
system-view
2. Enter OSPFv3 view.
ospfv3 [*process-id* | **vpn-instance** *vpn-instance-name*] *
3. Enable OSPFv3 link tag inheritance.
link-tag inherit enable
By default, OSPFv3 link tag inheritance is disabled.

Command reference

link-tag inherit enable

Use **link-tag inherit enable** to enable OSPFv3 link tag inheritance.

Use **undo link-tag inherit enable** to disable OSPFv3 link tag inheritance.

Syntax

```
link-tag inherit enable
undo link-tag inherit enable
```

Default

OSPFv3 link tag inheritance is disabled.

Views

OSPFv3 view

Predefined user roles

network-admin

Usage guidelines

After you execute this command on a device, the device inherits the OSPFv3 link tag in received LSPs. If the device receives multiple link tags from a device, it selects and inherits only one tag. Then, the device configures the link tag as the route tag for all routes that pass the link.

This feature is usually used in Source Address Validation Architecture (SAVA) scenarios to ensure the consistency of IPv6 SAVA entries on multiple gateway devices as follows:

1. Set an OSPFv3 link tag for the customer-side interface on the gateway device that advertises IPv6 SAVA entries.
2. Enable OSPFv3 link tag inheritance on the gateway device that receives IPv6 SAVA entries. The device generates an IPv6 SAVA entry based on a route prefix if the following two tags are the same:
 - The route tag configured by executing the **ipv6 sava import remote-route-tag** command.
 - The OSPFv3 link tag that the route inherits.

Examples

In OSPFv3 process 1, enable OSPFv3 link tag inheritance.

```
<Sysname> system-view
[Sysname] ospfv3 1
[Sysname-ospfv3-1] link-tag inherit enable
```

Related commands

```
ipv6 sava import remote-route-tag
ospfv3 link-tag
```

ospfv3 link-tag

Use **ospfv3 link-tag** to set an OSPFv3 link tag for an interface.

Use **undo ospfv3 link-tag** to restore the default.

Syntax

```
ospfv3 link-tag tag [ instance instance-id ]
undo ospfv3 link-tag [ instance instance-id ]
```

Default

No OSPFv3 link tag is configured for an interface.

Views

Interface view

Predefined user roles

network-admin

Parameters

tag: Specifies an OSPFv3 link tag in the range of 1 to 4294967295.

instance *instance-id*: Specifies an instance by its ID in the range of 0 to 255. The default instance ID is 0.

Usage guidelines

Execute this command together with the **link-tag inherit enable** command to filter routes based on OSPFv3 link tags.

If you execute this command multiple times on an interface, the most recent configuration takes effect.

Examples

```
# Set the OSPFv3 link tag for interface VLAN-interface10 to 222333.
```

```
<Sysname> system-view
```

```
[Sysname] interface vlan-interface 10
```

```
[Sysname-Vlan-interface10] ospfv3 link-tag 222333
```

Related commands

```
link-tag inherit enable
```

New feature: Displaying a message on enablement of the FTP or Telnet server

Displaying a message on enablement of the FTP or Telnet server

When you enable the FTP or Telnet server, the device displays a message indicating that FTP or Telnet might be insecure because it transmits all data in plaintext form.

Command reference

ftp server enable

Use **ftp server enable** to enable the FTP server.

Use **undo ftp server enable** to disable the FTP server.

Syntax

```
ftp server enable
```

```
undo ftp server enable
```

Default

The FTP server is disabled.

Views

System view

Predefined user roles

network-admin

Usage guidelines

FTP might be insecure because it transmits data in plaintext form.

Examples

```
# Enable the FTP server.
```

```
Sysname> system-view
[Sysname] ftp server enable
FTP is insecure because it transmits data in plaintext form.
```

telnet server enable

Use **telnet server enable** to enable the Telnet server.

Use **undo telnet server enable** to disable the Telnet server.

Syntax

```
telnet server enable
undo telnet server enable
```

Default

The Telnet server is disabled.

Views

System view

Predefined user roles

network-admin

Usage guidelines

This command is not supported in FIPS mode.

Users can Telnet to the device only when the Telnet server is enabled.

Telnet might be insecure because it transmits all data in plaintext form.

Examples

```
# Enable the Telnet server.
<Sysname> system-view
[Sysname] telnet server enable
Telnet is insecure because it transmits data in plaintext form.
```

New feature: Enabling OSPF to advertise the maximum link cost to neighbors

Enabling OSPF to advertise the maximum link cost to neighbors

About this task

In an OSPF network, when a link recovers from failures or the state of an interface changes, OSPF will re-establish neighbor relationships and perform route convergence. During the route convergence process, routing loops and traffic loss might occur because the convergence speeds of the nodes are different. To address this issue, enable OSPF to advertise the maximum link cost to neighbors within the specified period, so the traffic forwarding path remains unchanged. After the specified period, OSPF advertises the original link cost to neighbors and performs optimal route selection again.

Procedure

4. Enter system view.
system-view
5. Enter interface view.
interface *interface-type* *interface-number*
6. Enable OSPF to advertise the maximum link cost to neighbors within the specified period.
ospf peer hold-max-cost duration *time*
By default, OSPF does not advertise the maximum link cost to neighbors during a route convergence.

Command reference

ospf peer hold-max-cost duration

Use **ospf peer hold-max-cost duration** to enable OSPF to advertise the maximum link cost to neighbors within the specified period.

Use **undo ospf peer hold-max-cost duration** to restore the default.

Syntax

```
ospf peer hold-max-cost duration time  
undo ospf peer hold-max-cost duration
```

Default

OSPF does not advertise the maximum link cost to neighbors during a route convergence.

Views

Interface view

Predefined user roles

network-admin

Parameters

time: Specifies the time period during which OSPF advertises the maximum link cost to neighbors. The value range for this argument is 100 to 1000000 milliseconds.

Usage guidelines

In an OSPF network, when a link recovers from failures or the state of an interface changes, OSPF will re-establish neighbor relationships and perform route convergence. During the route convergence process, routing loops and traffic loss might occur because the convergence speeds of the nodes are different. To address this issue, enable OSPF to advertise the maximum link cost to neighbors within the specified period, so the traffic forwarding path remains unchanged. After the specified period, OSPF advertises the original link cost to neighbors and performs optimal route selection again.

The timer specified by the *time* argument starts and OSPF begins to advertise the maximum link cost (65535) to a neighbor when the OSPF neighbor enters Full state.

Examples

On interface VLAN-interface 10, enable OSPF to advertise the maximum link cost to neighbors within 1000 milliseconds.

```
<Sysname> system-view  
[Sysname] interface vlan-interface 10
```

```
[Sysname-Vlan-interface10] ospf peer hold-max-cost duration 1000
```

New feature: Enabling IS-IS to advertise the maximum link cost to neighbors

Enabling IS-IS to advertise the maximum link cost to neighbors

About this task

In an IS-IS network, when a link recovers from failures or the state of an interface changes, IS-IS will re-establish neighbor relationships and perform route convergence. During the route convergence process, routing loops and traffic loss might occur because the convergence speeds of the nodes are different. To address this issue, enable IS-IS to advertise the maximum link cost to neighbors within the specified period, so the traffic forwarding path remains unchanged. After the specified period, IS-IS advertises the original link cost to neighbors and performs optimal route selection again.

Procedure

1. Enter system view.
system-view
2. Enter interface view.
interface *interface-type* *interface-number*
3. Enable IS-IS to advertise the maximum link cost to neighbors within the specified period.
isis peer hold-max-cost duration *time*
By default, IS-IS does not advertise the maximum link cost to neighbors during a route convergence.

Command reference

isis peer hold-max-cost duration

Use **isis peer hold-max-cost duration** to enable IS-IS to advertise the maximum link cost to neighbors within the specified period.

Use **undo isis peer hold-max-cost duration** to restore the default.

Syntax

```
isis peer hold-max-cost duration time  
undo isis peer hold-max-cost duration
```

Default

IS-IS does not advertise the maximum link cost to neighbors during a route convergence.

Views

Interface view

Predefined user roles

network-admin

Parameters

time: Specifies the period during which IS-IS advertises the maximum link cost to neighbors, in the range of 100 to 1000000 milliseconds.

Usage guidelines

In an IS-IS network, when a link recovers from failures or the state of an interface changes, IS-IS will re-establish neighbor relationships and perform route convergence. During the route convergence process, routing loops and traffic loss might occur because the convergence speeds of the nodes are different. To address this issue, enable IS-IS to advertise the maximum link cost to neighbors within the specified period, so the traffic forwarding path remains unchanged. After the specified period, IS-IS advertises the original link cost to neighbors and performs optimal route selection again.

For a P2P neighbor, the timer specified by the *time* argument starts after the neighbor becomes up. For a broadcast neighbor, the timer specified by the *time* argument starts after the neighbor becomes up and the DIS is elected.

When the cost type is **wide**, the maximum link cost is 16777214. When the cost type is **Narrow**, the maximum link cost is 63.

This command applies to both IPv4 IS-IS and IPv6 IS-IS.

Examples

On interface VLAN-interface10, enable IS-IS to advertise the maximum link cost to neighbors within 1000 milliseconds.

```
<Sysname> system-view
[Sysname] interface vlan-interface 10
[Sysname-Vlan-interface10] isis peer hold-max-cost duration 1000
```

New feature: Enabling OSPFv3 to advertise the maximum link cost to neighbors

Enabling OSPFv3 to advertise the maximum link cost to neighbors

About this task

In an OSPFv3 network, when a link recovers from failures or the state of an interface changes, OSPFv3 will re-establish neighbor relationships and perform route convergence. During the route convergence process, routing loops and traffic loss might occur because the convergence speeds of the nodes are different. To address this issue, enable OSPFv3 to advertise the maximum link cost to neighbors within the specified period, so the traffic forwarding path remains unchanged. After the specified period, OSPFv3 advertises the original link cost to neighbors and performs optimal route selection again.

Procedure

1. Enter system view.
system-view
2. Enter interface view.
interface *interface-type* *interface-number*
3. Enable OSPFv3 to advertise the maximum link cost to neighbors within the specified period.
ospfv3 peer hold-max-cost duration *timer* [**instance** *instance-id*]

By default, OSPFv3 does not advertise the maximum link cost to neighbors during a route convergence.

Command reference

ospfv3 peer hold-max-cost duration

Use **ospfv3 peer hold-max-cost duration** to enable OSPFv3 to advertise the maximum link cost to neighbors within the specified period.

Use **undo ospfv3 peer hold-max-cost duration** to restore the default.

Syntax

```
ospfv3 peer hold-max-cost duration time [instance instance-id]  
undo ospfv3 peer hold-max-cost duration [instance instance-id]
```

Default

OSPFv3 does not advertise the maximum link cost to neighbors during a route convergence.

Views

Interface view

Predefined user roles

network-admin

Parameters

time: Specifies the time period during which OSPFv3 advertises the maximum link cost to neighbors in milliseconds. The value range for this argument is 100 to 1000000.

instance-id: Specifies an instance by its ID. The value range for this argument is 0 to 255. The default value is 0.

Usage guidelines

In an OSPFv3 network, when a link recovers from failures or the state of an interface changes, OSPFv3 will re-establish neighbor relationships and perform route convergence. During the route convergence process, routing loops and traffic loss might occur because the convergence speeds of the nodes are different. To address this issue, enable OSPFv3 to advertise the maximum link cost (65535) to neighbors within the specified period, so the traffic forwarding path remains unchanged. After the specified period, OSPFv3 advertises the original link cost to neighbors and performs optimal route selection again.

The timer specified by the *time* argument starts and OSPFv3 begins to advertise the maximum link cost (65535) to a neighbor when the neighbor enters Full state.

Examples

On interface VLAN-interface10, enable OSPFv3 to advertise the maximum link cost to neighbors within 1000 milliseconds.

```
<Sysname> system-view  
[Sysname] interface vlan-interface 10  
[Sysname-Vlan-interface10] ospfv3 peer hold-max-cost duration 1000
```

New feature: Enabling BGP to advertise a default route to a peer or peer group

Enabling BGP to advertise a default route to a peer or peer group

7. Enter system view.
system-view
8. Enter BGP instance view.
bgp *as-number* [**instance** *instance-name*]
9. Enter BGP EVPN address family view.
address-family **l2vpn** **evpn**
10. Enable BGP to advertise a default route to a peer or peer group.
peer { *group-name* | *ipv4-address* [*mask-length*] | *ipv6-address* [*prefix-length*] } **default-route-advertise** { **ipv4** | **ipv6** }
vpn-instance *vpn-instance-name*
By default, BGP does not advertise a default route to any peers or peer groups.

Command reference

Modified command: peer default-route-advertise

Use **peer default-route-advertise** to advertise a default route to a peer or peer group.

Use **undo peer default-route-advertise** to disable default route advertisement to a peer or peer group.

Old syntax

In BGP IPv4 multicast address family view:

```
peer { group-name | ipv4-address [ mask-length ] } default-route-advertise
[ route-policy route-policy-name ]

undo peer { group-name | ipv4-address [ mask-length ] }
default-route-advertise
```

In BGP VPNv4 address family view:

```
peer { group-name | ipv4-address [ mask-length ] } default-route-advertise
vpn-instance vpn-instance-name

undo peer { group-name | ipv4-address [ mask-length ] }
default-route-advertise vpn-instance vpn-instance-name
```

In BGP IPv4 unicast address family view/BGP-VPN IPv4 unicast address family view/BGP IPv6 unicast address family view/BGP-VPN IPv6 unicast address family view:

```
peer { group-name | ipv4-address [ mask-length ] | ipv6-address
[ prefix-length ] } default-route-advertise [ route-policy
route-policy-name ]

undo peer { group-name | ipv4-address [ mask-length ] | ipv6-address
[ prefix-length ] } default-route-advertise
```

In BGP IPv6 multicast address family view:

```
peer { group-name | ipv6-address [ prefix-length ] }
default-route-advertise [ route-policy route-policy-name ]
```

```
undo peer { group-name | ipv6-address [ prefix-length ] }  
default-route-advertise
```

In BGP IPv4 RT filter address family view:

```
peer { group-name | ipv4-address [ mask-length ] } default-route-advertise  
[ route-policy route-policy-name ]
```

```
undo peer { group-name | ipv4-address [ mask-length ] }  
default-route-advertise
```

New syntax

In BGP IPv4 multicast address family view:

```
peer { group-name | ipv4-address [ mask-length ] } default-route-advertise  
[ route-policy route-policy-name ]
```

```
undo peer { group-name | ipv4-address [ mask-length ] }  
default-route-advertise
```

In BGP VPNv4 address family view:

```
peer { group-name | ipv4-address [ mask-length ] } default-route-advertise  
vpn-instance vpn-instance-name
```

```
undo peer { group-name | ipv4-address [ mask-length ] }  
default-route-advertise vpn-instance vpn-instance-name
```

In BGP IPv4 unicast address family view/BGP-VPN IPv4 unicast address family view/BGP IPv6 unicast address family view/BGP-VPN IPv6 unicast address family view:

```
peer { group-name | ipv4-address [ mask-length ] | ipv6-address  
[ prefix-length ] } default-route-advertise [ route-policy  
route-policy-name ]
```

```
undo peer { group-name | ipv4-address [ mask-length ] | ipv6-address  
[ prefix-length ] } default-route-advertise
```

In BGP IPv6 multicast address family view:

```
peer { group-name | ipv6-address [ prefix-length ] }  
default-route-advertise [ route-policy route-policy-name ]
```

```
undo peer { group-name | ipv6-address [ prefix-length ] }  
default-route-advertise
```

In BGP IPv4 RT filter address family view:

```
peer { group-name | ipv4-address [ mask-length ] } default-route-advertise  
[ route-policy route-policy-name ]
```

```
undo peer { group-name | ipv4-address [ mask-length ] }  
default-route-advertise
```

In BGP EVPN address family view:

```
peer { group-name | ipv4-address [ mask-length ] | ipv6-address  
[ prefix-length ] } default-route-advertise { ipv4 | ipv6 } vpn-instance  
vpn-instance-name
```

```
undo peer { group-name | ipv4-address [ mask-length ] | ipv6-address  
[ prefix-length ] } default-route-advertise { ipv4 | ipv6 } vpn-instance  
vpn-instance-name
```

Old views

BGP IPv4 unicast address family view

BGP-VPN IPv4 unicast address family view

- BGP VPNv4 address family view
- BGP IPv6 unicast address family view
- BGP-VPN IPv6 unicast address family view
- BGP IPv4 multicast address family view
- BGP IPv6 multicast address family view
- BGP IPv4 RT filter address family view

New views

- BGP IPv4 unicast address family view
- BGP-VPN IPv4 unicast address family view
- BGP VPNv4 address family view
- BGP IPv6 unicast address family view
- BGP-VPN IPv6 unicast address family view
- BGP IPv4 multicast address family view
- BGP IPv6 multicast address family view
- BGP IPv4 RT filter address family view
- BGP EVPN address family view

Change description

Before modification: The `peer default-route-advertise` command was not supported in BGP EVPN address family view.

After modification: The `peer default-route-advertise` command is supported in BGP EVPN address family view.

New feature: SAVI

Configuring SAVI

About SAVI

Source Address Validation Improvement (SAVI) checks the validity of the source addresses of global unicast IPv6 packets. It implements the validity check by using the ND snooping, DHCPv6 snooping, ND attack detection, and IP source guard features. SAVI checks only global unicast addresses and forwards the packets that pass the validity check. Packets sourced from an invalid address are dropped.

SAVI application scenarios

DHCPv6-only

The hosts connected to the SAVI-enabled device obtain addresses only through DHCPv6. DHCPv6 messages, ND messages (RA and RR messages excluded), and IPv6 data packets are checked based on DHCPv6 snooping entries and static IPv6 source guard binding entries.

SLAAC-only

The hosts connected to the SAVI-enabled device obtain addresses only through Stateless Address Autoconfiguration (SLAAC). In this scenario, SAVI drops all DHCPv6 messages. Only ND messages

and IPv6 data packets are checked based on DHCPv6 snooping entries and static IPv6 source guard binding entries.

DHCPv6+SLAAC

The hosts connected to the SAVI-enabled device obtain addresses through DHCPv6 and SLAAC. In this scenario, SAVI checks all DHCPv6 messages, ND messages, and IPv6 data packets based on DHCPv6 snooping entries, ND snooping entries, and static IPv6 source guard binding entries.

SAVI tasks at a glance

To configure SAVI, perform the following tasks:

1. [Enabling SAVI](#)
2. [Configuring IPv6 source guard](#)
3. [Configuring DHCPv6 snooping](#)
4. [Configuring ND parameters](#)
5. (Optional.) [Setting the entry deletion delay](#)

Enabling SAVI

1. Enter system view.

```
system-view
```

2. Enable SAVI.

```
ipv6 savi strict
```

By default, SAVI is disabled.

Configuring IPv6 source guard

1. Enable IPv6 source guard on an interface.
2. (Optional.) Configure static IPv6SG bindings.

For more information about IPv6 source guard configuration, see *Security Configuration Guide*.

Configuring DHCPv6 snooping

Restrictions and guidelines

Enable only DHCPv6 snooping for the SLAAC-only scenario.

Procedure

1. Enable DHCPv6 snooping.
2. Specify DHCPv6 snooping trusted ports.
3. Enable recording client information in DHCPv6 snooping entries.

For more information about DHCPv6 snooping configuration, see *Layer 3—IP Services Configuration Guide*.

Configuring ND parameters

Restrictions and guidelines

Enable only ND attack detection for the DHCPv6-only scenario.

Procedure

1. Enable ND snooping for global unicast addresses.
For more information about ND snooping, see IPv6 basics in *Layer 3—IP Services Configuration Guide*.
2. Enable ND attack detection.
For more information about ND attack detection, see *Security Configuration Guide*.
3. Specify ND trusted ports.
For more information about ND trusted ports, see *Security Configuration Guide*.

Setting the entry deletion delay

About this task

The entry deletion delay is the period of time that the device waits before deleting the DHCPv6 snooping entries and ND snooping entries for a down port.

Procedure

1. Enter system view.
system-view
2. Set the entry deletion delay.
ipv6 savi down-delay *delay-time*
By default, the entry deletion delay is 30 seconds.

Enabling packet spoofing logging and filtering entry logging

About this task

Packet spoofing logging enables the device to generate log messages for the spoofed packets detected by SAVI.

Filtering entries are effective bindings used for filtering IPv6 packets by the source IPv6 address. Filtering entry logging enables the device to generate log messages for filtering entries. A log message contains the IPv6 address, MAC address, VLAN, and interface of a filtering entry.

The device sends packet spoofing and filtering entry log messages to the information center. With the information center, you can set log message filtering and output rules, including output destinations. For more information about using the information center, see *Network Management and Monitoring Configuration Guide*.

Procedure

1. Enter system view.
system-view
2. Enable packet spoofing logging.
ipv6 savi log enable spoofing-packet [**interval** *interval* | **total-number** *number*] *
By default, packet spoofing logging is disabled.
3. Enable filtering entry logging.
ipv6 savi log enable filter-entry
By default, filtering entry logging is disabled.

SAVI configuration examples

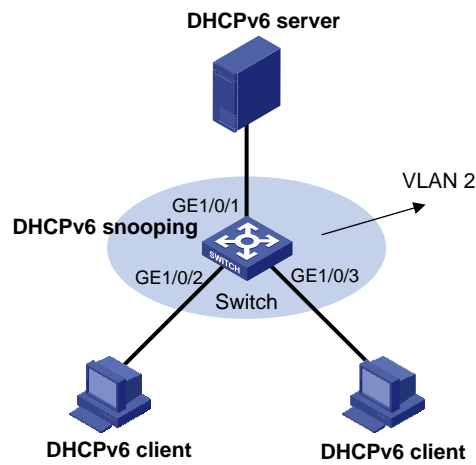
Example: Configuring DHCPv6-only SAVI

Network configuration

As shown in [Figure 2](#), configure SAVI on the switch to meet the following requirements:

- Clients obtain IPv6 addresses only through DHCPv6.
- SAVI checks the source addresses of DHCPv6 messages, ND messages (RA and RR messages excluded), and IPv6 data packets on Ten-GigabitEthernet1/0/2 and Ten-GigabitEthernet1/0/3.

Figure 2 Network diagram



Procedure

Enable SAVI.

```
<Switch> system-view
[Switch] ipv6 savi strict
```

Assign Ten-GigabitEthernet1/0/1 through Ten-GigabitEthernet1/0/3 to VLAN 2.

```
[Switch] vlan 2
[Switch-vlan2] port ten-gigabitethernet1/0/1 ten-gigabitethernet1/0/2
ten-gigabitethernet1/0/3
[Switch-vlan2] quit
```

Enable DHCPv6 snooping.

```
[Switch] ipv6 dhcp snooping enable
```

Configure Ten-GigabitEthernet1/0/1 as a DHCPv6 snooping trusted port.

```
[Switch] interface ten-gigabitethernet1/0/1
[Switch-Ten-GigabitEthernet1/0/1] ipv6 dhcp snooping trust
[Switch-Ten-GigabitEthernet1/0/1] quit
```

Enable recording DHCPv6 snooping entries on Ten-GigabitEthernet1/0/2 and Ten-GigabitEthernet1/0/3.

```
[Switch] interface ten-gigabitethernet1/0/2
[Switch-Ten-GigabitEthernet1/0/2] ipv6 dhcp snooping binding record
[Switch-Ten-GigabitEthernet1/0/2] quit
[Switch] interface ten-gigabitethernet1/0/3
[Switch-Ten-GigabitEthernet1/0/3] ipv6 dhcp snooping binding record
```

```
[Switch-Ten-GigabitEthernet1/0/3] quit
# Enable ND attack detection.
[Switch] vlan 2
[Switch-vlan2] ipv6 nd detection enable
[Switch-vlan2] quit
# Enable IPv6 source guard on Ten-GigabitEthernet1/0/2 and Ten-GigabitEthernet1/0/3.
[Switch] interface ten-gigabitethernet1/0/2
[Switch-Ten-GigabitEthernet1/0/2] ipv6 verify source ip-address mac-address
[Switch-Ten-GigabitEthernet1/0/2] quit
[Switch] interface ten-gigabitethernet1/0/3
[Switch-Ten-GigabitEthernet1/0/3] ipv6 verify source ip-address mac-address
[Switch-Ten-GigabitEthernet1/0/3] quit
```

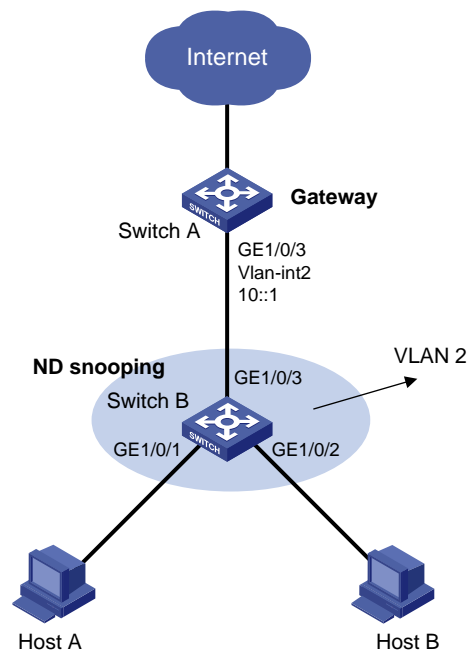
Example: Configuring SLAAC-only SAVI

Network configuration

As shown in [Figure 2](#), configure SAVI on Switch B to meet the following requirements:

- Hosts obtain IPv6 addresses only through SLAAC.
- DHCPv6 messages are dropped on Ten-GigabitEthernet1/0/1 through Ten-GigabitEthernet1/0/3 in VLAN 2.
- SAVI checks the source addresses of ND messages and IPv6 data packets on Ten-GigabitEthernet1/0/1 and Ten-GigabitEthernet1/0/2.

Figure 3 Network diagram



Procedure

```
# Enable SAVI.
<SwitchB> system-view
[SwitchB] ipv6 savi strict
# Assign Ten-GigabitEthernet1/0/1 through Ten-GigabitEthernet1/0/3 to VLAN 2.
```

```

[SwitchB] vlan 2
[SwitchB-vlan2] port ten-gigabitethernet1/0/1 ten-gigabitethernet1/0/2
ten-gigabitethernet1/0/3
[SwitchB-vlan2] quit

# Enable ND snooping for global unicast addresses in VLAN 2.
[SwitchB] vlan 2
[SwitchB-vlan2] ipv6 nd snooping enable global

# Enable ND attack detection for VLAN 2.
[SwitchB-vlan2] ipv6 nd detection enable
[SwitchB-vlan2] quit

# Enable DHCPv6 snooping.
[SwitchB] ipv6 dhcp snooping enable

# Configure Ten-GigabitEthernet1/0/3 as an ND trusted port.
[SwitchB] interface ten-gigabitethernet1/0/3
[SwitchB-Ten-GigabitEthernet1/0/3] ipv6 nd detection trust
[SwitchB-Ten-GigabitEthernet1/0/3] quit

# Enable IPv6 source guard on Ten-GigabitEthernet1/0/1 and Ten-GigabitEthernet1/0/2.
[SwitchB] interface ten-gigabitethernet1/0/1
[SwitchB-Ten-GigabitEthernet1/0/1] ipv6 verify source ip-address mac-address
[SwitchB-Ten-GigabitEthernet1/0/1] quit
[SwitchB] interface ten-gigabitethernet1/0/2
[SwitchB-Ten-GigabitEthernet1/0/2] ipv6 verify source ip-address mac-address
[SwitchB-Ten-GigabitEthernet1/0/2] quit

```

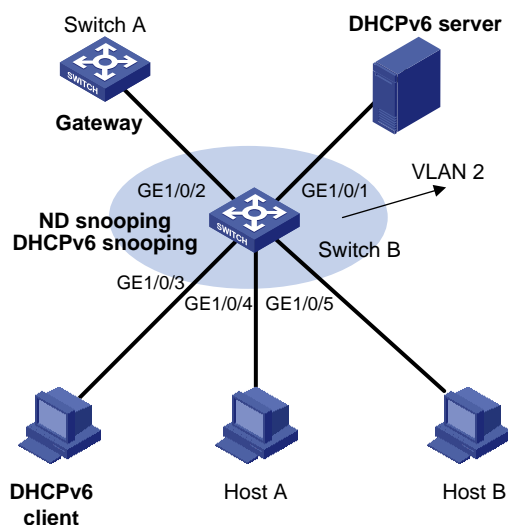
Example: Configuring DHCPv6+SLAAC SAVI

Network configuration

As shown in [Figure 3](#), configure SAVI on Switch B to meet the following requirements:

- Hosts obtain IP addresses through DHCPv6 or SLAAC.
- SAVI checks the source addresses of DHCPv6 messages, ND messages, and IPv6 data packets on Ten-GigabitEthernet1/0/3 through Ten-GigabitEthernet1/0/5.

Figure 4 Network diagram



Procedure

Enable SAVI.

```
<SwitchB> system-view
[SwitchB] ipv6 savi strict
```

Assign Ten-GigabitEthernet1/0/1 through Ten-GigabitEthernet1/0/5 to VLAN 2.

```
[SwitchB] vlan 2
[SwitchB-vlan2] port ten-gigabitethernet1/0/1 ten-gigabitethernet1/0/2
ten-gigabitethernet1/0/3 ten-gigabitethernet1/0/4 ten-gigabitethernet1/0/5
```

Enable DHCPv6 snooping.

```
[SwitchB] ipv6 dhcp snooping enable
```

Enable recording DHCPv6 snooping entries on Ten-GigabitEthernet1/0/3 through Ten-GigabitEthernet1/0/5.

```
[SwitchB] interface ten-gigabitethernet1/0/3
[SwitchB-Ten-GigabitEthernet1/0/3] ipv6 dhcp snooping binding record
[SwitchB-Ten-GigabitEthernet1/0/3] quit
[SwitchB] interface ten-gigabitethernet1/0/4
[SwitchB-Ten-GigabitEthernet1/0/4] ipv6 dhcp snooping binding record
[SwitchB-Ten-GigabitEthernet1/0/4] quit
[SwitchB] interface ten-gigabitethernet1/0/5
[SwitchB-Ten-GigabitEthernet1/0/5] ipv6 dhcp snooping binding record
[SwitchB-Ten-GigabitEthernet1/0/5] quit
```

Configure Ten-GigabitEthernet1/0/1 as a DHCPv6 snooping trusted port.

```
[SwitchB] interface ten-gigabitethernet1/0/1
[SwitchB-Ten-GigabitEthernet1/0/1] ipv6 dhcp snooping trust
[SwitchB-Ten-GigabitEthernet1/0/1] quit
```

Enable ND snooping for global unicast addresses in VLAN 2.

```
[SwitchB] vlan 2
[SwitchB-vlan2] ipv6 nd snooping enable global
```

Enable ND attack detection for VLAN 2.

```
[SwitchB-vlan2] ipv6 nd detection enable
[SwitchB-vlan2] quit
```

Configure Ten-GigabitEthernet1/0/2 as an ND trusted port.

```
[SwitchB] interface ten-gigabitethernet1/0/2
[SwitchB-Ten-GigabitEthernet1/0/2] ipv6 nd detection trust
[SwitchB-Ten-GigabitEthernet1/0/2] quit
```

Enable IPv6 source guard on Ten-GigabitEthernet1/0/3 through Ten-GigabitEthernet1/0/5.

```
[SwitchB] interface ten-gigabitethernet1/0/3
[SwitchB-Ten-GigabitEthernet1/0/3] ipv6 verify source ip-address mac-address
[SwitchB-Ten-GigabitEthernet1/0/3] quit
[SwitchB] interface ten-gigabitethernet1/0/4
[SwitchB-Ten-GigabitEthernet1/0/4] ipv6 verify source ip-address mac-address
[SwitchB-Ten-GigabitEthernet1/0/4] quit
[SwitchB] interface ten-gigabitethernet1/0/5
[SwitchB-Ten-GigabitEthernet1/0/5] ipv6 verify source ip-address mac-address
```

Command reference

ipv6 savi down-delay

Use **ipv6 savi down-delay** to set the entry deletion delay.

Use **undo ipv6 savi down-delay** to restore the default.

Syntax

```
ipv6 savi down-delay delay-time  
undo ipv6 savi down-delay
```

Default

The entry deletion delay is 30 seconds.

Views

System view

Predefined user roles

network-admin

Parameters

delay-time: Specifies the entry deletion delay in the range of 0 to 21474836 seconds.

Usage guidelines

The entry deletion delay is the period of time that the device waits before deleting the DHCPv6 snooping entries and ND snooping entries for a down port.

Examples

```
# Set the entry deletion delay to 100 seconds.  
<Sysname> system-view  
[Sysname] ipv6 savi down-delay 100
```

ipv6 savi log enable

Use **ipv6 savi log enable** to enable packet spoofing logging or filtering entry logging.

undo ipv6 savi log enable to disable packet spoofing logging or filtering entry logging.

Syntax

```
ipv6 savi log enable { spoofing-packet [ interval interval | total-number  
number ] * | filter-entry }  
undo ipv6 savi log enable { spoofing-packet | filter-entry }
```

Default

Packet spoofing logging and filtering entry logging are disabled.

Views

System view

Predefined user roles

network-admin

Parameters

spoofing-packet [**interval** *interval* | **total-number** *number*]: Enables packet spoofing logging.

- **interval** *interval*: Sets the log output interval in seconds. The value of the *interval* argument can be 0 or in the range of 5 to 3600. The default value is 60 seconds. If you set this parameter to 0, the device outputs a log message immediately after it is generated.
- **total-number** *number*: Sets the maximum number of log messages that can be output per interval. The value range for the *number* argument is 1 to 128, and the default value is 128.

filter-entry: Enables filtering entry logging.

Usage guidelines

Packet spoofing logging enables the device to generate log messages for the spoofed packets detected by SAVI.

A device can output a maximum of 128 packet spoofing log messages. If this limit is crossed, the device drops excess log messages. To ensure device performance, set the log output interval and maximum number of log messages output per interval appropriately.

Examples

```
# Enable packet spoofing logging.
<Sysname> system-view
[Sysname] ipv6 savi log enable spoofing-packet
```

ipv6 savi strict

Use **ipv6 savi strict** to enable Source Address Validation Improvement (SAVI).

Use **undo ipv6 savi strict** to disable SAVI.

Syntax

```
ipv6 savi strict
undo ipv6 savi strict
```

Default

SAVI is disabled.

Views

System view

Predefined user roles

network-admin

Examples

```
# Enable SAVI.
<Sysname> system-view
[Sysname] ipv6 savi strict
```

Related commands

```
ipv6 verify source
```

New feature: Setting the maximum number of IPv6SG bindings on an interface

Setting the maximum number of IPv6SG bindings on an interface

About this task

You can set the maximum number of IPv6SG bindings on an interface to limit the total number of static and dynamic IPv6SG bindings on the interface. If the upper limit is reached, no more IPv6SG bindings can be added on the interface.

Procedure

1. Enter system view.
system-view
2. Enter Layer 2 Ethernet interface view.
interface *interface-type interface-number*
3. Set the maximum number of IPv6SG bindings on the interface.
ipv6 verify source max-entries *number*

By default, the maximum number of IPv6SG bindings is not limited on an interface.

Command reference

ipv6 verify source max-entries

Use **ipv6 verify source max-entries** to set the maximum number of IPv6SG bindings on an interface.

Use **undo ipv6 verify source max-entries** to restore the default.

Syntax

```
ipv6 verify source max-entries number  
undo ipv6 verify source max-entries
```

Default

The maximum number of IPv6SG bindings is not limited on an interface.

Views

Layer 2 Ethernet interface view

Predefined user roles

network-admin

Parameters

number: Specifies the maximum number of IPv6SG bindings, in the range of 0 to 4294967295.

Usage guidelines

This command limits the total number of static and dynamic IPv6SG bindings on an interface.

If the maximum number of IPv6SG bindings is reached on an interface, no more IPv6SG bindings can be added on the interface.

If you set the maximum number of IPv6SG bindings to be smaller than the number of existing IPv6SG bindings, the configuration can be successful. The existing IPv6SG bindings are not affected.

The interface does not learn IPv6SG bindings until the number of existing IPv6SG bindings is smaller than the maximum number.

The number of effective IPv6SG bindings on an interface is limited by hardware resources.

Examples

```
# Set the maximum number of IPv6SG bindings to 200 on GigabitEthernet 1/0/1.
<Sysname> system-view
[Sysname] interface gigabitethernet 1/0/1
[Sysname-GigabitEthernet1/0/1] ipv6 verify source max-entries 200
```

Related commands

ipv6 source binding (interface view)

New feature: Setting the interval at which the SNMP module examines the system configuration for changes

Setting the interval at which the SNMP module examines the system configuration for changes

About this task

This task enables the SNMP module to examine the system configuration for changes at the specified intervals and generate a trap and a log if any change is found.

Procedure

1. Enter system view.
system-view
2. Set the interval at which the SNMP module examines the system configuration for changes.
snmp-agent configuration-examine interval *interval*
By default, the SNMP module examines the system configuration for changes at intervals of 600 seconds.

Command reference

snmp-agent configuration-examine interval

Use **snmp-agent configuration-examine interval** to set the interval at which the SNMP module examines the system configuration for changes.

Use **undo snmp-agent configuration-examine interval** to restore the default.

Syntax

```
snmp-agent configuration-examine interval interval
undo snmp-agent configuration-examine interval
```

Default

The SNMP module examines the system configuration for changes at intervals of 600 seconds.

Views

System view

Predefined user roles

network-admin

Parameters

interval: Specifies the interval at which the SNMP module examines the system configuration for changes, in seconds. The value is in the range of 1 to 86400.

Usage guidelines

This command enables the SNMP module to examine the system configuration for changes at the specified intervals and generate a log and a trap if any change is found.

Examples

```
# Set the interval at which the SNMP module examines the system configuration for changes to 600 seconds.
<sysname> system-view
[sysname] snmp-agent configuration-examine interval 600
```

New feature: Displaying SNMP notifications drop records

Displaying SNMP notifications drop records

When an SNMP notification is dropped from the SNMP trap queue, information about the notification is recorded in the SNMP trap buffer. You can execute the **display snmp-agent trapbuffer drop** command in any view to display the SNMP notifications drop records.

Command reference

display snmp-agent trapbuffer drop

Use **display snmp-agent trapbuffer drop** to display SNMP notifications drop records.

Syntax

```
display snmp-agent trapbuffer drop
```

Views

Any view

Predefined user roles

network-admin
network-operator

Usage guidelines

When an SNMP notification is dropped from the SNMP trap queue, information about the notification is recorded in the SNMP trap buffer.

Examples

```
# Display SNMP notifications drop records.
<Sysname> display snmp-agent trapbuffer drop
Current messages:1
```

Wed Dec 14 10:49:52:656 2019 Notification hh3cCfgManEventlog(1.3.6.1.4.1.25506.2.4.2.1) dropped.

Current messages in the command output indicates the total number of SNMP notifications drop records in the SNMP trap buffer.

Related commands

`reset snmp-agent trapbuffer`

New feature: Displaying SNMP notifications sending records

Displaying SNMP notifications sending records

After an SNMP notification is sent, information about the notification is recorded in the SNMP trap buffer. The information includes the content, destination IP address, and sending result of the notification. You can execute the **display snmp-agent trapbuffer send** command in any view to display SNMP notifications sending records.

Command reference

display snmp-agent trapbuffer send

Use **display snmp-agent trapbuffer send** to display SNMP notifications sending records.

Syntax

display snmp-agent trapbuffer send

Views

Any view

Predefined user roles

network-admin

network-operator

Usage guidelines

After an SNMP notification is sent, information about the notification is recorded in the SNMP trap buffer. The information includes the content, destination IP address, and sending result of the notification.

Examples

Display SNMP notifications sending records.

```
<Sysname> display snmp-agent trapbuffer send
```

Current messages:2

Dec 14 10:49:52:656 2019 Notification hh3cCfgManEventlog(1.3.6.1.4.1.25506.2.4.2.1) with hh3cCfgLogSrcCmd(1.3.6.1.4.1.25506.2.4.1.1.7.1.3.101)=2;hh3cCfgLogSrcData(1.3.6.1.4.1.25506.2.4.1.1.7.1.4.101)=4;hh3cCfgLogDesData(1.3.6.1.4.1.25506.2.4.1.1.7.1.5.101)=2 sent to 192.168.111.77 successfully.

Dec 14 10:49:52:659 2019 Notification hh3cCfgManEventlog(1.3.6.1.4.1.25506.2.4.2.1) with hh3cCfgLogSrcCmd(1.3.6.1.4.1.25506.2.4.1.1.7.1.3.101)=2;hh3cCfgLogSrcData(1.3.6.1.4.1.25506.2.4.1.1.7.1.4.101)=4;hh3cCfgLogDesData(1.3.6.1.4.1.25506.2.4.1.1.7.1.5.101)=2 failed to be sent to 192.168.111.77.

Current messages in the command output indicates the total number of SNMP notifications sending records in the SNMP trap buffer.

New feature: Clearing all records from the SNMP trap buffer

Clearing all records from the SNMP trap buffer

You can execute the `reset snmp-agent trapbuffer` command in user view to clear all records from the SNMP trap buffer.

Command reference

reset snmp-agent trapbuffer

Use `reset snmp-agent trapbuffer` to clear all records from the SNMP trap buffer.

Syntax

```
reset snmp-agent trapbuffer
```

Views

User view

Predefined user roles

network-admin

Examples

```
# Clear all records from the SNMP trap buffer.
```

```
<Sysname> reset snmp-agent trapbuffer
```

Related commands

```
display snmp-agent trapbuffer drop
```

```
display snmp-agent trapbuffer send
```

New feature: PWDCTL messages

PWDCTL_UPDATETIME

| | |
|---------------------------|---|
| Message text | Last login time updated after clock update. |
| Variable fields | N/A |
| Severity level | 6 |
| Example | PWDCTL/6/PWDCTL_UPDATETIME: Last login time updated after clock update. |
| Explanation | This message is sent when the last login time is updated. |
| Recommended action | No action is required. |

New feature: LIPC messages

PORT_CHANGE

| | |
|---------------------------|---|
| Message text | Last login time updated after clock update. |
| Variable fields | N/A |
| Severity level | 6 |
| Example | PWDCTL/6/PWDCTL_UPDATETIME: Last login time updated after clock update. |
| Explanation | This message is sent when the last login time is updated. |
| Recommended action | No action is required. |

New feature: Applying an SSL server policy to the NETCONF over SOAP over HTTPS service

Applying an SSL server policy to the NETCONF over SOAP over HTTPS service

About this task

You can apply an SSL server policy to the NETCONF over SOAP over HTTPS service. The NETCONF over SOAP over HTTPS service will use the SSL server policy to enhance service security. For more information about SSL server policies, see SSL configuration in *Security Configuration Guide*.

After NETCONF over SOAP over HTTPS is enabled, changes to the applied SSL server policy do not affect established NETCONF over SOAP over HTTPS sessions. The changes affect only NETCONF over SOAP over HTTPS sessions established after the changes are made.

Procedure

1. Enter system view.
system-view
2. Apply an SSL server policy to the NETCONF over SOAP over HTTPS service.
netconf soap https ssl-server-policy *policy-name*
By default, no SSL server policy is applied to the NETCONF over SOAP over HTTPS service.

Command reference

netconf soap https ssl-server-policy

Use **netconf soap https ssl-server-policy** to apply an SSL server policy to the NETCONF over SOAP over HTTPS service.

Use **undo netconf soap https ssl-server-policy** to restore the default.

Syntax

netconf soap https ssl-server-policy *policy-name*

```
undo netconf soap https ssl-server-policy
```

Default

No SSL server policy is applied to the NETCONF over SOAP over HTTPS service.

Views

System view

Predefined user roles

network-admin

Parameters

policy-name: Specifies an SSL server policy name, a string of 1 to 31 characters.

Usage guidelines

The NETCONF over SOAP over HTTPS service will use the SSL server policy to enhance service security. For more information about SSL server policies, see SSL configuration in *Security Configuration Guide*.

You can configure this command only when NETCONF over SOAP over HTTPS is disabled.

This command takes effect after you enable NETCONF over SOAP over HTTPS.

If you execute this command multiple times, the most recent configuration takes effect.

After NETCONF over SOAP over HTTPS is enabled, changes to the applied SSL server policy do not affect established NETCONF over SOAP over HTTPS sessions. The changes affect only NETCONF over SOAP over HTTPS sessions established after the changes are made.

Examples

```
# Apply SSL server policy myssl to the NETCONF over SOAP over HTTPS service.
```

```
<Sysname> system-view
```

```
[Sysname] netconf soap https ssl-server-policy myssl
```

New feature: Policing ND protocol packet to the control plane

Configuring ND protocol packet policing

Restrictions and guidelines

By default, ND protocol packets to the control plane are policed to 1000 pps. The default policed rate value can be viewed in the **display qos policy control-plane pre-defined** command output. You can modify the policed rate value by applying a QoS policy to the control plane.

Procedure

1. Enter system view.
system-view
2. Create a traffic class and enter traffic class view.
traffic classifier *classifier-name* [**operator** { **and** | **or** }]
3. Configure a match criterion.
if-match [**not**] *match-criteria*
By default, no match criteria are configured.
4. Return to system view.

- quit**
5. Create a traffic behavior and enter traffic behavior view.
traffic behavior *behavior-name*
 6. Configure a traffic policing action.

```
car cir [ pps ] committed-information-rate [ cbs committed-burst-size ]
[ ebs excess-burst-size ] [ green action | red action | yellow action ]
*
```

```
car cir [ pps ] committed-information-rate [ cbs committed-burst-size ]
pir [ pps ] peak-information-rate [ ebs excess-burst-size ] [ green
action | red action | yellow action ] *
```

By default, no traffic policing action is configured.
 7. Return to system view.
quit
 8. Create a QoS policy and enter QoS policy view.
qos policy *policy-name*
 9. Associate the class with the traffic behavior in the QoS policy.

```
classifier classifier-name behavior behavior-name [ insert-before
before-classifier-name ]
```

By default, a class is not associated with any behavior.
 10. Return to system view.
quit
 11. Enter control plane view.
control-plane slot *slot-number*
 12. Apply the QoS policy to the control plane.

By default, no QoS policy is applied to a control plane.

Command reference

Modified command: display qos policy control-plane pre-defined

Syntax

```
display qos policy control-plane pre-defined [ slot slot-number ]
```

Views

Any view

Change description

Before modification: No information is displayed for the ND protocol in the command output.

After modification: The information about the ND protocol appears in the command output.

Modified command: if-match

Syntax

```
if-match match-criteria
```

Views

Traffic class view

Change description

Before modification: The `nd` value is not supported in the `protocol-name` argument of the `if-match control-plane protocol protocol-name` command.

After modification: The `nd` value was added to the `protocol-name` argument of the `if-match control-plane protocol protocol-name` command.

New feature: RA message transmission over VXLAN tunnels between distributed EVPN gateways

Enabling a distributed EVPN gateway to send RA messages over VXLAN tunnels

About this task

By default, a distributed EVPN gateway drops the RS messages received from VXLAN tunnels and periodically advertises RA messages only to the local site. As a result, a distributed EVPN gateway does not send RA messages over VXLAN tunnels, and remote gateways cannot update information about the gateway based on RA messages. To resolve the issue, perform this task to enable distributed EVPN gateways to reply to remote RS messages with RA messages and periodically advertise RA messages over VXLAN tunnels.

Procedure

1. Enter system view.
system-view
 2. Enter VSI interface view.
interface vsi-interface vsi-interface-id
 3. Enable the VSI interface to send RA messages over VXLAN tunnels.
ipv6 nd ra tunnel-broadcast enable
- By default, a VSI interface does not send RA messages over VXLAN tunnels.

Command reference

ipv6 nd ra tunnel-broadcast enable

Use **ipv6 nd ra tunnel-broadcast enable** to enable a VSI interface to send RA messages over VXLAN tunnels.

Use **undo ipv6 nd ra tunnel-broadcast enable** to disable a VSI interface from sending RA messages over VXLAN tunnels.

Syntax

```
ipv6 nd ra tunnel-broadcast enable
undo ipv6 nd ra tunnel-broadcast enable
```

Default

A VSI interface does not send RA messages over VXLAN tunnels.

Views

VSI interface view

Predefined user roles

network-admin

Usage guidelines

By default, a distributed EVPN gateway drops the RS messages received from VXLAN tunnels and periodically advertises RA messages only to the local site. As a result, a distributed EVPN gateway does not send RA messages over VXLAN tunnels, and remote gateways cannot update information about the gateway based on RA messages. To resolve the issue, use this command to enable distributed EVPN gateways to reply to remote RS messages with RA messages and periodically advertise RA messages over VXLAN tunnels.

Examples

```
# Enable VSI-interface 100 to send RA messages over VXLAN tunnels.
<Sysname> system-view
[Sysname] interface vsi-interface 100
[Sysname-Vsi-interface100] ipv6 nd ra tunnel-broadcast enable
```

New feature: MAC move logging for VSIs

Displaying MAC move records for VSIs

You can execute the **display l2vpn mac-address mac-move** command in any view to display the MAC move events that have occurred on VSIs since startup of the device.

Command reference

display l2vpn mac-address mac-move

Use **display l2vpn mac-address mac-move** to display MAC move records for VSIs.

Syntax

```
display l2vpn mac-address mac-move [ slot slot-number ]
```

Views

Any view

Predefined user roles

network-admin
network-operator

Parameters

slot *slot-number*: Specifies an IRF member device by its member ID. If you do not specify a member device, the command displays MAC move records for the master device.

Usage guidelines

When a MAC address frequently moves between two interfaces, Layer 2 loops might occur in the network. To discover and locate loops, you can view the MAC address move records.

This command displays the MAC move events that have occurred on VSIs since startup of the device. The MAC move records with the same MAC address, VSI name, source port, and current port are considered to be one record.

An IRF member device can generate a maximum of 200 MAC address move records.

Examples

Display MAC move records for all VSIs.

```
<Sysname> display l2vpn mac-address mac-move
```

```
MAC address : 0000-4900-03e7      VSI name   : 3
```

```
Current port: HGE1/0/1 srv14      Source port: HGE1/0/1 srv31
```

```
Last time   : 2019-02-19 20:44:15 Move count : 1
```

```
MAC address : 0000-4900-03e8      VSI name   : 3
```

```
Current port: HGE1/0/1 srv14      Source port: HGE1/0/1 srv31
```

```
Last time   : 2019-02-19 20:44:16 Move count : 1
```

```
--- 2 MAC address moving records found ---
```

Table 4 Command output

| Field | Description |
|--------------|--|
| Current port | Interface to which the MAC address was moved. |
| Source port | Interface from which the MAC address was moved. |
| Last time | Last time when the MAC address was moved. |
| Move count | Number of times that the MAC address has moved from the source port to the current port. |

```
reset l2vpn mac-address mac-move
```

Use **reset l2vpn mac-address mac-move** to clear MAC move records for all VSIs.

Syntax

```
reset l2vpn mac-address mac-move
```

Views

User view

Predefined user roles

network-admin

network-operator

Examples

Clear MAC move records for all VSIs.

```
<Sysname> reset l2vpn mac-address mac-move
```

New feature: MAC learning limits for VSIs

Setting the MAC learning limit on a VSI

Restrictions and guidelines

The MAC learning limit controls the number of MAC addresses learned from local ACs.

Procedure

1. Enter system view.

- system-view**
- 2. Enter VSI view.
vsi *vsi-name*
- 3. Set a limit for the VSI's MAC address table.
mac-table limit *mac-limit*
By default, no limit is set for a VSI's MAC address table.
- 4. Enable the VSI to drop source-unknown unicast frames if the MAC address table is full.
mac-table limit drop-unknown
By default, a VSI forwards source-unknown unicast frames without learning the source MAC address if the MAC address table is full.

Command reference

mac-table limit

Use **mac-table limit** to set a limit for the MAC address table of a VSI.

Use **undo mac-table limit** to restore the default.

Syntax

```
mac-table limit mac-limit
undo mac-table limit
```

Default

No limit is set for a VSI's MAC address table.

Views

VSI view

Predefined user roles

network-admin

Parameters

mac-limit: Specifies the maximum number of MAC addresses that can be learned by the VSI. The value range for this argument is 0 to 65535.

Usage guidelines

This command helps limit the MAC address table size of VSIs to save resources.

Examples

```
# Configure VSI vpn1 to learn a maximum of 1024 MAC address entries.
<Sysname> system-view
[Sysname] vsi vpn1
[Sysname-vsi-vpn1] mac-table limit 1024
```

Related commands

```
display l2vpn vsi
```

mac-table limit drop-unknown

Use **mac-table limit drop-unknown** to enable a VSI to drop source-unknown unicast frames if the MAC address table is full.

Use `undo mac-table limit drop-unknown` to restore the default.

Syntax

```
mac-table limit drop-unknown
```

```
undo mac-table limit drop-unknown
```

Default

A VSI forwards source-unknown unicast frames without learning the source MAC address if the MAC address table is full.

Views

VSI view

Predefined user roles

network-admin

Examples

Configure VSI **vpn1** to learn a maximum of 1024 MAC address entries, and enable it to drop source-unknown unicast frames if the MAC address table is full.

```
<Sysname> system-view
```

```
[Sysname] vsi vpn1
```

```
[Sysname-vsi-vpn1] mac-table limit 1024
```

```
[Sysname-vsi-vpn1] mac-table limit drop-unknown
```

Related commands

```
display l2vpn vsi
```

New feature: DRNI in EVPN-DCI

Configuring DRNI in EVPN-DCI

About this task

To set up a DR system with two EDs, configure a virtual VTEP address on the EDs. The EDs will use the virtual VTEP address to set up VXLAN tunnels or VXLAN-DCI tunnels with VTEPs or remote EDs.

Restrictions and guidelines

Do not execute the **evpn drni local** command if you have configured DRNI in EVPN-DCI.

When you attach a user site to a DR system, attach it to both DR interfaces in a DR group. Do not configure single-homed ACs on the member EDs.

Procedure

1. Enter system view.

```
system-view
```

2. Specify the virtual VTEP address.

```
evpn drni group { virtual-vtep-ipv4 / virtual-vtep-ipv6 }
```

By default, DRNI in EVPN-DCI is not configured.

3. Enter BGP instance view.

```
bgp as-number [ instance instance-name ]
```

4. Enter BGP EVPN address family view.

address-family l2vpn evpn

5. Enable the device to replace the next hop in advertised BGP EVPN routes with the virtual VTEP address.

nexthop evpn-drni group-address

The default settings are as follows:

- When advertising BGP EVPN routes to an EBGp peer or peer group, the device replaces the next hop with the IP address of the source interface used to establish BGP sessions.
- When advertising EBGp routes to an IBGP peer or peer group, the device does not modify the next hop.

Command reference

nexthop evpn-drni group-address

Use **nexthop evpn-drni group-address** to enable the device to replace the next hop in advertised BGP EVPN routes with the virtual VTEP address.

Use **nexthop evpn-drni group-address** to restore the default.

Syntax

nexthop evpn-drni group-address

undo nexthop evpn-drni group-address

Default

When advertising BGP EVPN routes to an EBGp peer or peer group, the device replaces the next hop with the IP address of the source interface used to establish BGP sessions. When advertising EBGp routes to an IBGP peer or peer group, the device does not modify the next hop.

Views

BGP EVPN address family view

Predefined user roles

network-admin

Usage guidelines

When you use DRNI on EDs in an EVPN-DCI network, you must enable the EDs to replace the next hop in advertised BGP EVPN routes with the virtual VTEP address.

Examples

Enable the device to replace the next hop in advertised BGP EVPN routes with the virtual VTEP address.

```
<Sysname> system-view
```

```
[Sysname] bgp 100
```

```
[Sysname-bgp-default] address-family l2vpn evpn
```

```
[Sysname-bgp-default-evpn] nexthop evpn-drni group-address
```

New feature: L2VPN statistics collection interval

Setting the L2VPN statistics collection interval

About this task

Perform this task to set the interval for collecting VSI, AC, and VXLAN tunnel packet statistics.

Procedure

1. Enter system view.
system-view
2. Set the L2VPN statistics collection interval
l2vpn statistics interval *interval*
The default L2VPN statistics collection interval is 15 minutes.

Command reference

l2vpn statistics interval

Use **l2vpn statistics interval** to set the L2VPN statistics collection interval.

Use **undo l2vpn statistics interval** to restore the default.

Syntax

```
l2vpn statistics interval interval  
undo l2vpn statistics interval
```

Default

The L2VPN statistics collection interval is 15 minutes.

Views

System view

Predefined user roles

network-admin

Parameters

Interval: Sets the L2VPN statistics collection interval in seconds, in the range of 30 to 65535.

Examples

Set the L2VPN statistics collection interval to 30 seconds.

```
<Sysname> system-view  
[Sysname] l2vpn statistics interval 30
```

New feature: Unsolicited NA learning

Enabling unsolicited NA learning

About this task

On some networks, a server multicasts NA messages to two peer devices for link backup. The peer devices cannot learn ND entry for the server from these NA messages by default. If no ND learning is triggered by data exchange between the server and peer devices, the peer devices learn the entry for the server only when the server unicasts messages to them.

This feature enables an interface to learn ND entries from unsolicited NA messages. The ND entries generated by using this method are in stale state.

Restrictions and guidelines

To ensure that the device learns ND entries from trusted NA messages, enable this feature only on a secure network.

This feature might cause the device to learn excessive ND entries that consume too many system resources. As a best practice, execute the **ipv6 neighbor stale-aging** command to set a smaller aging timer before you enable this feature. The smaller aging timer accelerates the aging of ND entries in stale state.

Procedure

1. Enter system view.
system-view
 2. Enter Layer 3 interface view.
interface *interface-type* *interface-number*
 3. Enable unsolicited NA learning.
ipv6 nd unsolicited-na-learning enable
- By default, unsolicited NA learning is disabled.

Command reference

ipv6 nd unsolicited-na-learning enable

Use **ipv6 nd unsolicited-na-learning enable** to enable unsolicited NA learning.

Use **undo ipv6 nd unsolicited-na-learning enable** to disable unsolicited NA learning.

Syntax

```
ipv6 nd unsolicited-na-learning enable
undo ipv6 nd unsolicited-na-learning enable
```

Default

Unsolicited NA learning is disabled.

Views

Layer 3 interface view

Predefined user roles

network-admin
network-operator

Usage guidelines

To ensure that the device learns ND entries from trusted NA messages, enable this feature only on a secure network.

This feature might cause the device to learn excessive ND entries that consume too many system resources. As a best practice, execute the **ipv6 neighbor stale-aging** command to set a smaller aging timer before you enable this feature. The smaller aging timer accelerates the aging of ND entries in stale state.

Examples

```
# Enable unsolicited NA learning on VLAN-interface 2.
```

```
<Sysname> system-view
```

```
[Sysname] interface vlan-interface 2
```

```
[Sysname-Vlan-interface2] ipv6 nd unsolicited-na-learning enable
```

Related commands

```
ipv6 neighbor stale-aging
```

New feature: IPv6 enhancement for OpenFlow flow entries

Enabling IPv6 enhancement for OpenFlow flow entries

About this task

If two extensibility flow tables are configured on the device, ACL hardware entries are shared by flow entries matching IPv6 packets and other flow entries. This might cause flow entry issuing failure. To resolve this issue, use this command so the device will use independent ACL hardware resources to store flow entries in extensibility flow tables used to match IPv6 packets.

Procedure

1. Enter system view.
system-view
2. Enable IPv6 enhancement for OpenFlow flow entries.
openflow flow-table ipv6-enhanced
By default, IPv6 enhancement is disabled for OpenFlow flow entries.

Command reference

openflow flow-table ipv6-enhanced

Use **openflow flow-table ipv6-enhanced** to enable IPv6 enhancement for OpenFlow flow entries.

Use **undo openflow flow-table ipv6-enhanced** to disable IPv6 enhancement for OpenFlow flow entries.

Syntax

```
openflow flow-table ipv6-enhanced
```

```
undo openflow flow-table ipv6-enhanced
```

Default

IPv6 enhancement for OpenFlow flow entries is disabled.

Views

System view

Predefined user roles

network-admin

Usage guidelines

If two extensibility flow tables are configured on the device, ACL hardware entries are shared by flow entries matching IPv6 packets and other flow entries. This might cause flow entry issuing failure. To resolve this issue, use this command so the device will use independent ACL hardware resources to store flow entries in extensibility flow tables used to match IPv6 packets.

Execution of this command reactivates all active OpenFlow instances and clears all existing flow entries.

Examples

```
# Enable IPv6 enhancement for OpenFlow flow entries.
```

```
<Sysname> system-view
```

```
[Sysname] openflow flow-table ipv6-enhanced
```

Related commands

active instance

flow-table

New feature: IPv6-matching QoS policy

Configuring an IPv6-matching QoS policy

About the IPv6-matching QoS policy

An IPv6-matching QoS policy is typically used for IPv6 packets matched by IPv6 ACLs.

You can apply a generic QoS policy used for IPv4 packets and an IPv6-matching QoS policy to the same direction of an interface.

Procedure

1. Enter system view.
system-view
2. Define a traffic class.
 - a. Create a traffic class and enter traffic class view.
traffic classifier *classifier-name* [**operator** { **and** | **or** }]
 - b. (Optional.) Configure a description for the traffic class.
description *text*
 - c. Configure a match criterion.
if-match *match-criteria*

By default, no match criterion is configured.
For configurable match criteria, see the **if-match** command in *ACL and QoS Command Reference*.
 - d. Return to system view.

- quit**
3. Define a traffic behavior.
 - a. Create a traffic behavior and enter traffic behavior view.
`traffic behavior behavior-name`
 - b. Configure an action.
 By default, no action is configured.
 - c. Return to system view.
`quit`
 4. Define a QoS policy.
 - a. Create an IPv6-matching QoS policy and enter QoS policy view.
`qos ipv6-matching policy policy-name`
 - b. Associate the class with the traffic behavior in the QoS policy.
`classifier classifier-name behavior behavior-name [mode { dcbx | qppb-manipulation }] | insert-before before-classifier-name]`
 By default, a traffic class is not associated with a traffic behavior.
 - c. Return to system view.
`quit`
 5. Apply the IPv6-matching QoS policy. Choose one option as needed:
 - o Execute the following commands in sequence to apply the IPv6-matching QoS policy to an interface:
`interface interface-type interface-number`
`qos apply ipv6-matching policy policy-name { inbound | outbound }`
`[share-mode]`
 - o Apply the IPv6-matching QoS policy globally:
`qos apply ipv6-matching policy policy-name global { inbound | outbound }`

By default, no QoS policy is applied.

Command reference

Modified command: qos policy

Old syntax

```
qos [ accounting | remark ] policy policy-name
undo qos [ accounting | remark ] policy policy-name
```

New syntax

```
qos [ ipv6-matching | { accounting | remark } ] policy policy-name
undo qos [ ipv6-matching | { accounting | remark } ] policy policy-name
```

Views

System view

Parameters

ipv6-matching: Specifies an IPv6-matching QoS policy.

Change description

Before modification: IPv6-matching QoS policies are not supported.

After modification: IPv6-matching QoS policies are supported.

Modified command: qos apply policy

Old syntax

```
qos apply [ accounting | remarking ] policy policy-name { inbound |  
outbound } [ share-mode ]  
  
undo qos apply [ accounting | remarking ] policy policy-name { inbound  
| outbound }
```

New syntax

```
qos apply [ ipv6-matching | { accounting | remarking } ] policy policy-name  
{ inbound | outbound } [ share-mode ]  
  
undo qos apply [ ipv6-matching | { accounting | remarking } ] policy  
policy-name { inbound | outbound }
```

Views

- Control plane view
- Ethernet service instance view
- Layer 2 Ethernet interface view
- Layer 2 aggregate interface view
- Layer 3 Ethernet interface view
- Layer 3 Ethernet subinterface view
- Layer 3 aggregate interface view
- S-channel interface view
- S-channel aggregate interface view
- Tunnel interface view
- VSI interface view
- VSI view

Parameters

ipv6-matching: Specifies an IPv6-matching QoS policy. This keyword is supported only in interface view.

Change description

Before modification: IPv6-matching QoS policies are not supported.

After modification: IPv6-matching QoS policies are supported.

Modified command: qos apply policy global

Old syntax

```
qos apply [ accounting | remarking ] policy policy-name global { inbound  
| outbound }  
  
undo qos apply [ accounting | remarking ] policy policy-name global  
{ inbound | outbound }
```

New syntax

```
qos apply [ ipv6-matching | { accounting | remarking } ] policy policy-name
global { inbound | outbound }

undo qos apply [ ipv6-matching | { accounting | remarking } ] policy
policy-name global { inbound | outbound }
```

Views

System view

Parameters

ipv6-matching: Specifies an IPv6-matching QoS policy.

Change description

Before modification: IPv6-matching QoS policies are not supported.

After modification: IPv6-matching QoS policies are supported.

Modified command: display qos policy

Old syntax

```
display qos policy user-defined [ accounting | remarking ] [ policy-name
[ classifier classifier-name ] ] [ slot slot-number ]
```

New syntax

```
display qos policy user-defined [ ipv6-matching | { accounting |
remarking } ] [ policy-name [ classifier classifier-name ] ] [ slot
slot-number ]
```

Views

Any view

Parameters

ipv6-matching: Specifies an IPv6-matching QoS policy.

Change description

Before modification: IPv6-matching QoS policies are not supported.

After modification: IPv6-matching QoS policies are supported.

Modified command: display qos policy interface

Old syntax

```
display qos [ accounting | remarking ] policy interface [ interface-type
interface-number ] [ slot slot-number ] [ inbound | outbound ]
```

New syntax

```
display qos [ ipv6-matching | { accounting | remarking } ] policy interface
[ interface-type interface-number ] [ slot slot-number ] [ inbound |
outbound ]
```

Views

Any view

Parameters

ipv6-matching: Specifies an IPv6-matching QoS policy.

Change description

Before modification: IPv6-matching QoS policies are not supported.

After modification: IPv6-matching QoS policies are supported.

Modified command: display qos policy global

Old syntax

```
display qos [ accounting | remarking ] policy global [ slot slot-number ]  
[ inbound | outbound ]
```

New syntax

```
display qos [ ipv6-matching | { accounting | remarking } ] policy global  
[ slot slot-number ] [ inbound | outbound ]
```

Views

Any view

Parameters

ipv6-matching: Specifies an IPv6-matching QoS policy.

Change description

Before modification: IPv6-matching QoS policies are not supported.

After modification: IPv6-matching QoS policies are supported.

Modified command: reset qos policy global

Old syntax

```
reset qos [ accounting | remarking ] policy global [ inbound | outbound ]
```

New syntax

```
reset qos [ ipv6-matching | { accounting | remarking } ] policy global  
[ inbound | outbound ]
```

Views

User view

Parameters

ipv6-matching: Specifies an IPv6-matching QoS policy.

Change description

Before modification: IPv6-matching QoS policies are not supported.

After modification: IPv6-matching QoS policies are supported.

New feature: Enabling BGP to not flush specific routes to the routing table

Enabling BGP to not flush specific routes to the routing table

About this task

Perform this task to enable BGP to not flush routes matching the specified routing policy to the routing table.

This feature applies to the following scenarios:

- In an MPLS L3VPN network, all the optimal private BGP routes will be flushed to the routing table of a PE by default. You can configure this feature to prevent unnecessary routes from being flushed to the IP routing table and improve the forwarding performance of the PE.
- In a route reflection network, configure this feature to not flush specific BGP routes to the routing table of the RR. Then, the RR only advertises and receives routes but does not forward service traffic, which saves the system resources of the RR.

Procedure (IPv4 unicast)

1. Enter system view.
system-view
2. Enter BGP IPv4 unicast address family view or BGP-VPN IPv4 unicast address family view.
 - Execute the following commands in sequence to enter BGP IPv4 unicast address family view:
bgp *as-number* [**instance** *instance-name*]
address-family ipv4 [**unicast**]
 - Execute the following commands in sequence to enter BGP-VPN IPv4 unicast address family view:
bgp *as-number* [**instance** *instance-name*]
ip vpn-instance *vpn-instance-name*
address-family ipv4 [**unicast**]
3. Enable BGP to not flush routes matching the specified routing policy to the routing table.
routing-table bgp-rib-only [**route-policy** *route-policy-name*]
By default, BGP flushes the optimal routes to the routing table.

Procedure (IPv6 unicast)

1. Enter system view.
system-view
2. Enter BGP IPv6 unicast address family view or BGP-VPN IPv6 unicast address family view.
 - Execute the following commands in sequence to enter BGP IPv6 unicast address family view:
bgp *as-number* [**instance** *instance-name*]
address-family ipv6 [**unicast**]
 - Execute the following commands in sequence to enter BGP-VPN IPv6 unicast address family view:
bgp *as-number* [**instance** *instance-name*]
ip vpn-instance *vpn-instance-name*
address-family ipv6 [**unicast**]

3. Enable BGP to not flush routes matching the specified routing policy to the routing table.

routing-table bgp-rib-only [**route-policy** *route-policy-name*]

By default, BGP flushes the optimal routes to the routing table.

Command reference

routing-table bgp-rib-only

Use **routing-table bgp-rib-only** to enable BGP to not flush routes matching the specified routing policy to the routing table.

Use **undo routing-table bgp-rib-only** to restore the default.

Syntax

routing-table bgp-rib-only [**route-policy** *route-policy-name*]

undo routing-table bgp-rib-only

Default

BGP flushes the optimal routes to the routing table.

Views

BGP IPv4 unicast address family view

BGP IPv6 unicast address family view

BGP-VPN IPv4 unicast address family view

BGP-VPN IPv6 unicast address family view

Predefined user roles

network-admin

Parameters

route-policy *route-policy-name*: Specifies a routing policy by its name, a case-sensitive string of 1 to 63 characters. If you do not specify a routing policy or the specified routing policy does not exist, BGP does not flush any routes to the routing table. The apply clauses in the specified routing policy do not take effect.

Usage guidelines

This command applies to the following scenarios:

- In an MPLS L3VPN network, all the optimal private BGP routes will be flushed to the routing table of a PE by default. You can execute this command to prevent unnecessary routes from being flushed to the IP routing table and improve the forwarding performance of the PE.
- In a route reflection network, execute this command to not flush specific BGP routes to the routing table of the RR. Then, the RR only advertises and receives routes but does not forward service traffic, which saves the system resources of the RR.

This command is mutually exclusive with the **advertise-rib-active** command.

Examples

In BGP IPv4 unicast address family view, enable BGP to not flush BGP IPv4 unicast routes matching routing policy **policy1** to the routing table.

```
<Sysname> system-view
```

```
[Sysname] bgp 100
```

```
[Sysname-bgp-default] address-family ipv4 unicast
```

```
[Sysname-bgp-default-ipv4] routing-table bgp-rib-only route-policy policy1
```

New feature: Enabling BGP to reset peer sessions gracefully

Enabling BGP to reset peer sessions gracefully

About this task

When the TCP connection goes down, the hold timer expires, or the support for a new address family is added, BGP tears down and then re-establishes the peer sessions, which will cause traffic interruption. To avoid traffic interruption in these cases, enable BGP to reset peer sessions gracefully.

Procedure

1. Enter system view.
system-view
2. Enter BGP instance view.
bgp *as-number* [**instance** *instance-name*]
3. Enable GR capability for BGP.
graceful-restart
By default, GR capability is disabled for BGP.
4. Enable BGP to reset peer sessions gracefully.
graceful-restart peer-reset [**all**]
By default, BGP does not reset peer sessions gracefully.

Command reference

graceful-restart peer-reset

Use **graceful-restart peer-reset** to enable BGP to reset peer sessions gracefully.

Use **undo graceful-restart peer-reset** to disable BGP from resetting peer sessions gracefully.

Syntax

```
graceful-restart peer-reset [ all ]  
undo graceful-restart peer-reset
```

Default

BGP does not reset peer sessions gracefully.

Views

BGP instance view

Predefined user roles

network-admin

Parameters

all: Enables BGP to reset peer sessions gracefully when the TCP connection goes down, the hold timer expires, or the support for a new address family is added. If you do not specify this keyword, BGP resets peer sessions gracefully only when the support for a new address family is added.

Usage guidelines

When the TCP connection goes down, the hold timer expires, or the support for a new address family is added, BGP tears down and then re-establishes the peer sessions, which will cause traffic interruption. To avoid traffic interruption in these cases, execute this command to enable BGP to reset peer sessions gracefully.

Examples

```
# Enable BGP to reset peer sessions gracefully.
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] graceful-restart
[Sysname-bgp-default] graceful-restart peer-reset
```

New feature: Setting the extra time to wait after the restart timer expires

Setting the extra time to wait after the restart timer expires

About this task

After the active/standby switchover or BGP restart completes, the GR helper marks the routes it learned from the GR restarter as stale routes. If the GR helper fails to establish a BGP session after both the GR timer and the extra timer to wait expire, the GR helper removes the stale routes.

Procedure

1. Enter system view.
system-view
2. Enter BGP instance view.
bgp *as-number* [**instance** *instance-name*]
3. Set the extra time to wait after the restart timer expires.
peer { *group-name* | *ipv4-address* [*mask-length*] | *ipv6-address* [*prefix-length*] } **graceful-restart timer restart extra** { *time* | **no-limit** }

By default, the extra time to wait after the restart timer expires is 0 seconds.

Command reference

peer graceful-restart timer restart extra

Use **peer graceful-restart timer restart extra** to set the extra time to wait after the restart timer expires.

Use **undo peer graceful-restart timer restart extra** to restore the default.

Syntax

```
peer { group-name | ipv4-address [ mask-length ] | ipv6-address [ prefix-length ] } graceful-restart timer restart extra { time | no-limit }

undo peer { group-name | ipv4-address [ mask-length ] | ipv6-address [ prefix-length ] } graceful-restart timer restart extra
```

Default

The extra time to wait after the restart timer expires is 0 seconds.

Views

BGP instance view

BGP-VPN instance view

Predefined user roles

network-admin

Parameters

group-name: Specifies a peer group by its name, a case-sensitive string of 1 to 47 characters. The peer group must have been created.

ipv4-address: Specifies a peer by its IPv4 address. The peer must have been created.

mask-length: Specifies a mask length in the range of 0 to 32. You can use the *ipv4-address* and *mask-length* arguments together to specify a subnet. If you specify a subnet, this command specifies all dynamic peers in the subnet.

ipv6-address: Specifies a peer by its IPv6 address. The peer must have been created.

prefix-length: Specifies a prefix length in the range of 0 to 128. You can use the *ipv6-address* and *prefix-length* arguments together to specify a subnet. If you specify a subnet, this command specifies all dynamic peers in the subnet.

time: Specifies the extra time to wait after the restart timer expires, in the range of 0 to 86400 seconds.

no-limit: Sets an unlimited time to wait for BGP session re-establishment.

Usage guidelines

After the active/standby switchover or BGP restart completes, the GR helper marks the routes it learned from the GR restarter as stale routes. If the GR helper fails to establish a BGP session after both the GR timer and the extra timer to wait expire, the GR helper removes the stale routes.

Examples

```
# Set the extra time to wait after the restart timer expires to 100 seconds.
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] peer 10.1.100.2 graceful-restart timer restart extra 100
```

New feature: Enabling password change prompt logging

Enabling password change prompt logging

About this task

Use this feature to enhance the protection of passwords for Telnet, SSH, HTTP, HTTPS, NETCONF over SSH, and NETCONF over SOAP users and improve the system security.

This feature enables the device to generate logs to prompt users to change their weak passwords at an interval of 24 hours and at a login basis.

A password is a weak password if it does not meet the following requirements:

- Password composition restriction configured by using the **password-control composition** command.
- Minimum password length restriction set by using the **password-control complexity** command.
- It cannot contain the username or the reverse letters of the username.

For a NETCONF over SSH or NETCONF over SOAP user, the device also generate a password change prompt log if any of the following conditions exists:

- The current password of the user is the default password or has expired.
- The user logs in to the device for the first time or uses a new password to log in after global password control is enabled.

The device will no longer generate password change prompt logs for a user when one of the following conditions exists:

- The password change prompt logging feature is disabled.
- The user has changed the password and the new password meets the password control requirements.
- The enabling status of a related password control feature has changed so the current password of the user meets the password control requirements.
- The password composition policy or the minimum password length has changed.

You can use the **display password-control** command to display password control configuration. For more information about password control commands, see password control commands in *Security Command Reference*.

Procedure

1. Enter system view.
system-view
2. Enable password change prompt logging.
local-server log change-password-prompt
By default, password change prompt logging is enabled.

Command reference

local-server log change-password-prompt

Use **local-server log change-password-prompt** to enable password change prompt logging.

Use **undo local-server log change-password-prompt** to disable password change prompt logging.

Syntax

```
local-server log change-password-prompt
undo local-server log change-password-prompt
```

Default

Password change prompt logging is enabled.

Views

System view

Predefined user roles

network-admin

Examples

```
# Enable password change prompt logging.  
<Sysname> system-view  
[Sysname] local-server log change-password-prompt
```

New feature: Enabling mandatory weak password change

Enabling mandatory weak password change

Login control with a weak password

The system checks for weak passwords for Telnet, SSH, HTTP, or HTTPS device management users. A password is a weak password if it does not meet the following requirements:

- Password composition restriction.
- Minimum password length restriction.
- It cannot contain the username or the reverse letters of the username.

By default, the system displays a message about a weak password but does not force the user to change the password. To improve the device security, you can enable the mandatory weak password change feature, which forces the users with weak passwords to change them. Then, the users can log in to the device only after their passwords meet the password restriction requirements.

First login

For some device models, the factory defaults contain a default username and password. If the device starts up with factory defaults, Telnet, SSH, HTTP, or HTTPS device management users must change the default password at first login before they can access the system.

Procedure

1. Enter system view.
system-view
2. Enable mandatory weak password change.
password-control change-password weak-password enable
By default, the mandatory weak password change feature is disabled.

Command reference

New command: password-control change-password weak-password enable

Use **password-control change-password weak-password enable** to enable mandatory weak password change.

Use **undo password-control change-password weak-password enable** to disable mandatory weak password change.

Syntax

```
password-control change-password weak-password enable  
undo password-control change-password weak-password enable
```

Default

The mandatory weak password change feature is disabled.

Views

System view

Predefined user roles

network-admin
network-operator

Usage guidelines

By default, a password must meet the following requirements:

- A password using the global composition policy must contain a minimum of two character types and a minimum of one character for each type.
- The global minimum password length is 10 characters.

You can use the **password-control composition** and **password-control length** commands to configure the global composition policy and minimum password length.

Examples

In system view, enable mandatory weak password change.

```
<Sysname> system-view
```

```
[Sysname] password-control change-password weak-password enable
```

Modified command: display password-control

Use **display password-control** to display password control configuration.

Syntax

```
display password-control [ super ]
```

Views

Any view

Predefined user roles

network-admin

Parameters

super: Displays the password control information for the super passwords. If you do not specify this keyword, the command displays the global password control configuration.

Usage guidelines

Before modification: The **Password change** field does not contain the enabling state of the mandatory weak password change feature.

After modification: The **Password change** field displays the enabling state of the mandatory weak password change feature.

Display the global password control configuration.

```
<Sysname> display password-control
```

```
Global password control configurations:
```

| | |
|-------------------|-----------------------------------|
| Password control: | Enabled (device management users) |
| | Enabled (network access users) |
| Password aging: | Enabled (90 days) |
| Password length: | Enabled (10 characters) |

| | |
|---------------------------------------|--|
| Password composition: | Enabled (1 types, 1 characters per type) |
| Password history: | Enabled (max history records:4) |
| Early notice on password expiration: | 7 days |
| User authentication timeout: | 600 seconds |
| Maximum login attempts: | 3 |
| Action for exceeding login attempts: | Lock user for 1 minutes |
| Password history was last reset: | 0 days ago (device management users) |
| | 0 days ago (network access users) |
| Minimum interval between two updates: | 24 hours |
| User account idle time: | 90 days |
| Logins with aged password: | 3 times in 30 days |
| Password complexity: | Disabled (username checking) |
| | Disabled (repeated characters checking) |
| Password change: | Enabled (first login) |
| | Enabled (mandatory weak password change) |

New feature: Configuring ND packet rate limit

Configuring ND packet rate limit

About this task

This feature limits the receiving rate of ND packets that are to be delivered to the CPU, preventing the CPU from being overwhelmed by ND packets.

Enabled with ND attack detection, the device redirects the received ND packets to the CPU for security check. If attackers send a large number of forged ND packets to the device, the CPU will be overloaded. As a result, other features might malfunction and the device might crash. To avoid this situation, configure ND packet rate limit on ND packet receiving interfaces. When the receiving rate of ND packets on the interfaces exceeds the rate limit, the packets that exceed the limit will be discarded.

When logging for ND packet rate limit is enabled, the device sends the highest threshold-crossed ND packet rate within the sending interval in a log message to the information center. You can configure the information center module to set the log output rules. For more information about information center, see *Network Management and Monitoring Configuration Guide*.

Restrictions and guidelines

As a best practice, configure this feature if the device is configured with ND attack detection or ND snooping or has detected ND flood attacks.

To prevent the device from sending log messages frequently, set the log message sending interval to a high value.

On a Layer 2 aggregate interface, if you enable ND packet rate limit and the logging for this feature, the device sends log messages upon limit violations on its member ports.

Procedure

1. Enter system view.
system-view
2. (Optional.) Enable logging for ND packet rate limit.
ipv6 nd rate-limit log enable

By default, logging for ND packet rate limit is disabled.

3. (Optional.) Set the log message sending interval.

```
ipv6 nd rate-limit log interval interval
```

By default, the device sends log messages every 60 seconds.

4. Enter interface view.

```
interface interface-type interface-number
```

Supported interfaces include Layer 2 Ethernet interfaces, Layer 2 aggregate interfaces, Layer 3 Ethernet interfaces, and Layer 3 aggregate interfaces.

5. Enable ND packet rate limit.

```
ipv6 nd rate-limit [ pps ]
```

By default, ND packet rate limit is enabled.

Command reference

ipv6 nd rate-limit

Use **ipv6 nd rate-limit** to enable ND packet rate limit.

Use **undo ipv6 nd rate-limit** to disable ND packet rate limit.

Syntax

```
ipv6 nd rate-limit [ pps ]
```

```
undo ipv6 nd rate-limit
```

Default

ND packet rate limit is enabled.

Views

Layer 2 Ethernet interface view

Layer 2 aggregate interface view

Layer 3 Ethernet interface view

Layer 3 aggregate interface view

Predefined user roles

network-admin

Parameters

pps: Specifies the upper limit for ND packet receiving rate, in pps. The value range for this argument is 5 to 2000, and the default is 2000. If you do not specify the limit, the default value applies. The default value varies by device model.

Usage guidelines

The rate limit limits the receiving rate of ND packets that are to be delivered to the CPU, preventing the CPU from being overwhelmed by ND packets. Packets that exceed the rate limit are dropped.

Examples

Enable ND packet rate limit on Layer 2 Ethernet interface GigabitEthernet 1/0/1, and set the rate limit to 50 pps.

```
<Sysname> system-view
```

```
[Sysname] interface gigabitethernet 1/0/1
```

```
[Sysname-GigabitEthernet1/0/1] ipv6 nd rate-limit 50
```

ipv6 nd rate-limit log enable

Use `ipv6 nd rate-limit log enable` to enable logging for ND packet rate limit.

Use `undo ipv6 nd rate-limit log enable` to disable logging for ND packet rate limit.

Syntax

```
ipv6 nd rate-limit log enable
undo ipv6 nd rate-limit log enable
```

Default

Logging for ND packet rate limit is disabled.

Views

System view

Predefined user roles

network-admin

Usage guidelines

When logging for ND packet rate limit is enabled, the device sends the highest threshold-crossed ND packet rate within the sending interval in a log message to the information center. You can configure the information center module to set the log output rules. For more information about information center, see *Network Management and Monitoring Configuration Guide*.

Examples

```
# Enable logging for ND packet rate limit.
<Sysname> system-view
[Sysname] ipv6 nd rate-limit log enable
```

Related commands

```
ipv6 nd rate-limit log interval
```

ipv6 nd rate-limit log interval

Use `ipv6 nd rate-limit log interval` to set the log message sending interval for ND packet rate limit.

Use `undo ipv6 nd rate-limit log interval` to restore the default.

Syntax

```
ipv6 nd rate-limit log interval interval
undo ipv6 nd rate-limit log interval
```

Default

The device sends log messages every 60 seconds when the ND packet receiving rate on an interface exceeds the limit.

Views

System view

Predefined user roles

network-admin

Parameters

interval: Specifies an interval in the range of 1 to 86400 seconds.

Usage guidelines

To change the default interval and activate it, you must enable ND packet rate limit and enable sending log messages for ND packet rate limit.

Examples

Configure the device to send log messages every 120 seconds when the ND packet receiving rate on an interface exceeds the limit.

```
<Sysname> system-view
[Sysname] ipv6 nd rate-limit log interval 120
```

Related commands

```
ipv6 nd rate-limit log enable
```

New feature: gRPC logging in dial-in mode

Enabling gRPC logging in dial-in mode

About this task

To log gRPC packet operations in dial-in mode, enable gRPC logging in dial-in mode.

This feature generates gRPC operation logs in dial-in mode and sends them to the information center.

With the information center, you can configure log destinations and output rules. For more information about the information center, see *Network Management and Monitoring Configuration Guide*.

Restrictions and guidelines

gRPC operation logging might degrade device performance if gRPC operations are frequent. As a best practice, use gRPC operation logging only when necessary and log only gRPC operations of interest if gRPC operation logging is enabled.

Procedure

1. Enter system view.
system-view
2. Enable gRPC logging in dial-in mode. Choose the options to configure as needed:
 - Enable gRPC logging for RPC operations in dial-in mode.
grpc log dial-in rpc { all | { cli | get }* }
By default, gRPC logging is disabled for RPC operations in dial-in mode.
 - Enable gRPC logging for gNMI operations in dial-in mode.
grpc log dial-in gnmi { all | { capabilities | get | set | subscribe }* }
By default, gRPC logging is enabled for gNMI Set operations and disabled for other gNMI operations in dial-in mode.

Command reference

grpc log dial-in gnmi

Use **grpc log dial-in gnmi** to enable gRPC logging for gNMI operations in dial-in mode.

Use **undo grpc log dial-in gnmi** to disable gRPC logging for gNMI operations in dial-in mode.

Syntax

```
grpc log dial-in gnmi { all | { capabilities | get | set | subscribe }* }  
undo grpc log dial-in gnmi { all | { capabilities | get | set | subscribe }* }
```

Default

In dial-in mode, gRPC logging is enabled for gNMI Set operations and disabled for other gNMI operations.

Views

System view

Predefined user roles

network-admin

Parameters

all: Specifies all gNMI operations.

capabilities: Specifies the gNMI Capabilities operations.

get: Specifies the gNMI Get operations.

set: Specifies the gNMI Set operations.

subscribe: Specifies the gNMI Subscribe operations.

Usage guidelines

To identify gRPC issues, enable gNMI operations logging.

This command generates gNMI operation logs in dial-in mode and sends them to the information center. With the information center, you can configure log destinations and output rules. For more information about the information center, see *Network Management and Monitoring Configuration Guide*.

If you execute this command multiple times, the most recent configuration takes effect.

Examples

```
# Enable gRPC logging for gNMI Get operations in dial-in mode.
```

```
<Sysname> system
```

```
[Sysname] grpc log dial-in gnmi get
```

grpc log dial-in rpc

Use **grpc log dial-in rpc** to enable gRPC logging for RPC operations in dial-in mode.

Use **undo grpc log dial-in rpc** to disable gRPC logging for RPC operations in dial-in mode.

Syntax

```
grpc log dial-in rpc { all | { cli | get }* }  
undo grpc log dial-in rpc { all | { cli | get }* }
```

Default

In dial-in mode, gRPC logging is disabled for RPC operations.

Views

System view

Predefined user roles

network-admin

Parameters

all: Specifies both RPC CLI and Get operations.

cli: Specifies the RPC CLI operations.

get: Specifies the RPC Get operations.

Usage guidelines

To identify gRPC issues, enable RPC operations logging.

This command generates RPC operation logs in dial-in mode and sends them to the information center. With the information center, you can configure log destinations and output rules. For more information about the information center, see *Network Management and Monitoring Configuration Guide*.

If you execute this command multiple times, the most recent configuration takes effect.

Examples

Enable gRPC logging for RPC Get operations in dial-in mode.

```
<Sysname> system
```

```
[Sysname] grpc log dial-in gnmi get
```

New feature: gRPC logging in dial-out mode

Enabling gRPC logging in dial-out mode

About this task

To log gRPC sampling in dial-out mode, enable gRPC logging in dial-out mode.

This feature generates gRPC sampling logs in dial-out mode and sends them to the information center.

With the information center, you can configure log destinations and output rules. For more information about the information center, see *Network Management and Monitoring Configuration Guide*.

Procedure

1. Enter system view.

```
system-view
```

2. Enable gRPC logging in dial-out mode.

```
grpc log dial-out { all | { event | sample }* }
```

By default, gRPC logging is disabled in dial-out mode.

Command reference

grpc log dial-out

Use **grpc log dial-out** to enable gRPC logging in dial-out mode.

Use **undo grpc log dial-out** to disable gRPC logging in dial-out mode.

Syntax

```
grpc log dial-out { all | { event | sample }* }  
undo grpc log dial-out { all | { event | sample }* }
```

Default

In dial-out mode, gRPC logging is disabled.

Views

System view

Predefined user roles

network-admin

Parameters

all: Specifies all sampling types.

event: Specifies event-triggered sampling.

sample: Specifies periodical sampling.

Usage guidelines

To identify gRPC issues, enable gRPC sampling logging in dial-out mode.

This command generates gRPC sampling logs in dial-out mode and sends them to the information center. With the information center, you can configure log destinations and output rules. For more information about the information center, see *Network Management and Monitoring Configuration Guide*.

If you execute this command multiple times, the most recent configuration takes effect.

Examples

```
# Enable gRPC logging for periodical sampling in dial-out mode.
```

```
<Sysname> system
```

```
[Sysname] grpc log dial-in gnmi get
```

New feature: Creating frame match criteria based on VXLAN IDs for dynamic ACs on the IPL of an EVPN DR system

Creating frame match criteria based on VXLAN IDs for dynamic ACs on the IPL of an EVPN DR system

About this task

By default, if an EVPN DR system uses an Ethernet aggregate link as the IPL, each DR member device creates a dynamic AC on the IPL when an AC is configured on a site-facing interface. The dynamic AC and the site-facing AC have the same frame match criterion and VSI mapping. If two site-facing ACs on different interfaces have the same frame match criterion but different VSI mappings, the dynamic ACs created for the site-facing ACs will conflict with each other. To prevent this issue, enable the DR member devices to create frame match criteria based on VXLAN IDs for the dynamic ACs on the IPL.

Procedure

1. Enter system view.

system-view

2. Enable the device to create frame match criteria based on VXLAN IDs for the dynamic ACs on the Ethernet aggregate link IPL.

l2vpn drni peer-link ac-match-rule vxlan-mapping

By default, on an EVPN DR system that uses an Ethernet aggregate link as the IPL, dynamic ACs on the IPL use frame match criteria that are identical to those of site-facing ACs.

If you do not execute this command, do not configure overlapping outer VLAN IDs for Ethernet service instances of different VSIs.

If you execute this command, do not create VXLANs with IDs larger than 16000000.

Command reference

l2vpn drni peer-link ac-match-rule vxlan-mapping

Use **l2vpn drni peer-link ac-match-rule vxlan-mapping** to enable the device to create frame match criteria based on VXLAN IDs for the dynamic ACs on the Ethernet aggregate link IPL.

Use **undo l2vpn drni peer-link ac-match-rule vxlan-mapping** to restore the default.

Syntax

l2vpn drni peer-link ac-match-rule vxlan-mapping

undo l2vpn drni peer-link ac-match-rule vxlan-mapping

Default

On an EVPN DR system that uses an Ethernet aggregate link as the IPL, dynamic ACs on the IPL use frame match criteria that are identical to those of site-facing ACs.

Views

System view

Predefined user roles

network-admin

Usage guidelines

By default, if an EVPN DR system uses an Ethernet aggregate link as the IPL, each DR member device creates a dynamic AC on the IPL when an AC is configured on a site-facing interface. The dynamic AC and the site-facing AC have the same frame match criterion and VSI mapping. If two site-facing ACs on different interfaces have the same frame match criterion but different VSI mappings, the dynamic ACs created for the site-facing ACs will conflict with each other. To prevent this issue, enable the DR member devices to create frame match criteria based on VXLAN IDs for the dynamic ACs on the IPL.

With this command configured, a DR member device creates dynamic ACs on the IPL and maps them to the VSIs of VXLANs after the VXLANs are created. The matching VLAN IDs in frame match criteria are calculated for the dynamic ACs according to the following rules:

- Outer VLAN ID = VXLAN ID / 4094 + 1.
- Inner VLAN ID = VXLAN ID % 4094 + 1.
- If the calculated outer VLAN ID of a dynamic AC is the PVID of the IPP, the device uses the calculated inner VLAN ID as the outer VLAN ID. The device does not add a matching inner VLAN ID to the frame match criterion of the dynamic AC.

After you execute this command, do not create VXLANs with IDs larger than 16000000.

This command deletes existing dynamic ACs from IPPs and takes effect on all VXLANs.

Examples

```
# Enable the device to create frame match criteria based on VXLAN IDs for the dynamic ACs on the Ethernet aggregate link IPL.
```

```
<Sysname> system-view
```

```
[Sysname] l2vpn drni peer-link ac-match-rule vxlan-mapping
```

New feature: Removing Layer 3 VPN route targets from BGP EVPN routes for VXLAN mapping

Removing Layer 3 VPN route targets from BGP EVPN routes for VXLAN mapping

About this task

In an EVPN-DCI network that is configured with the **mapping vni** command, use this feature to remove the route targets of the VPN instance associated with an L3 VXLAN ID from the BGP EVPN routes advertised to remote sites. With this feature, the remote sites do not create Layer 3 entries for the BGP EVPN routes.

Restrictions and guidelines

If both the original and re-originated routes are advertised to remote sites, the routes advertised to remote sites might carry different VXLAN IDs for the same entry. This issue causes forwarding failure.

Procedure

1. Enter system view.
system-view
2. Enter BGP instance view.
bgp as-number [instance instance-name]
3. Enter BGP EVPN address family view.
address-family l2vpn evpn
4. Remove Layer 3 VPN route targets from the BGP EVPN routes advertised to remote sites when VXLAN mapping is performed.
mapping-vni remove vpn-target

By default, the Layer 3 VPN route targets are carried in the BGP EVPN routes advertised to remote sites when VXLAN mapping is performed.

Command reference

mapping-vni remove vpn-target

Use **mapping-vni remove vpn-target** to remove Layer 3 VPN route targets from the BGP EVPN routes advertised to remote sites when VXLAN mapping is performed.

Use **undo mapping-vni remove vpn-target** to restore the default.

Syntax

```
mapping-vni remove vpn-target
```

```
undo mapping-vni remove vpn-target
```

Default

The Layer 3 VPN route targets are carried in the BGP EVPN routes advertised to remote sites when VXLAN mapping is performed.

Views

BGP EVPN address family view

Predefined user roles

network-admin

Usage guidelines

In an EVPN-DCI network that is configured with the **mapping vni** command, use this feature to remove the route targets of the VPN instance associated with an L3 VXLAN ID from the BGP EVPN routes advertised to remote sites. With this feature, the remote sites do not create Layer 3 entries for the BGP EVPN routes.

If both the original and re-originated routes are advertised to remote sites, the routes advertised to remote sites might carry different VXLAN IDs for the same entry. This issue causes forwarding failure.

Examples

In BGP EVPN address family view, remove Layer 3 VPN route targets from the BGP EVPN routes advertised to remote sites when VXLAN mapping is performed.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] address-family l2vpn evpn
[Sysname-bgp-default-evpn] mapping-vni remove vpn-target
```

Related commands

mapping vni

New feature: Enabling asymmetric IRB mode for EVPN VXLAN traffic forwarding

Configuring the traffic forwarding mode for EVPN VXLAN

Restrictions and guidelines

The asymmetric IRB mode is supported only on distributed EVPN gateways and takes effect only on Layer 3 traffic forwarded in the same VXLAN. In addition, the same VSI interface on different distributed EVPN gateways must have different IP addresses.

Procedure

1. Enter system view.
system-view
2. Configure the traffic forwarding mode for EVPN VXLAN. Choose one of the following options:
 - Enable asymmetric IRB mode.
evpn irb asymmetric
 - Enable symmetric IRB mode.
undo evpn irb asymmetric

By default, a distributed EVPN gateway forwards EVPN VXLAN traffic in symmetric IRB mode.

Command reference

evpn irb asymmetric

Use **evpn irb asymmetric** to enable asymmetric IRB mode for EVPN VXLAN traffic forwarding.

Use **undo evpn irb asymmetric** to restore the default.

Syntax

```
evpn irb asymmetric
undo evpn irb asymmetric
```

Default

A distributed EVPN gateway forwards EVPN VXLAN traffic in symmetric IRB mode.

Views

System view

Predefined user roles

network-admin

Usage guidelines

Use this command to change the traffic forwarding mode for EVPN VXLAN. EVPN VXLAN supports asymmetric IRB and symmetric IRB modes.

The asymmetric IRB mode is supported only on distributed EVPN gateways and takes effect only on Layer 3 traffic forwarded in the same VXLAN. In addition, the same VSI interface on different distributed EVPN gateways must have different IP addresses.

Examples

```
# Enable asymmetric IRB mode for EVPN VXLAN traffic forwarding.
<Sysname> system-view
[Sysname] evpn irb asymmetric
```

New feature: Setting the Gateway IP Address field for EVPN IP prefix advertisement routes

Setting the Gateway IP Address field for EVPN IP prefix advertisement routes

About this task

Perform this task in an EVPN network to modify the Gateway IP Address field for EVPN IP prefix advertisement routes.

Procedure

1. Enter system view.
system-view
2. Enter routing policy node view.
route-policy *route-policy-name* { **deny** | **permit** } **node** *node-number*
By default, BGP does not redistribute IGP routes.
3. Set the Gateway IP Address field for EVPN IP prefix advertisement routes.

IPv4:

```
apply gateway-ip { ip-address | origin-nexthop }
```

IPv6:

```
apply ipv6 gateway-ip { ipv6-address | origin-nexthop }
```

By default, the Gateway IP Address field is not set for EVPN IP prefix advertisement routes.

Command reference

apply gateway-ip

Use **apply gateway-ip** to set the Gateway IP Address field (IPv4) for EVPN IP prefix advertisement routes.

Use **undo apply gateway-ip** to restore the default.

Syntax

```
apply gateway-ip { ip-address | origin-nexthop }
```

```
undo apply gateway-ip
```

Default

The Gateway IP Address field is not set for EVPN IP prefix advertisement routes.

Views

Routing policy node view

Predefined user roles

network-admin

Parameters

ip-address: Specifies an IP address in dotted decimal notation.

origin-nexthop: Specifies the IP address of the route next hop.

Usage guidelines

Use this command in an EVPN network to modify the Gateway IP Address field for EVPN IP prefix advertisement routes.

Examples

```
# Set the Gateway IP Address field for EVPN IP prefix advertisement routes to the IP address of the route next hop.
```

```
<Sysname> system-view
```

```
[Sysname] route-policy policy1 permit node 10
```

```
[Sysname-route-policy-policy1-10] apply gateway-ip origin-nexthop
```

Related commands

```
apply ipv6 gateway-ip
```

apply ipv6 gateway-ip

Use **apply ipv6 gateway-ip** to set the Gateway IP Address field (IPv6) for EVPN IP prefix advertisement routes.

Use **undo apply ipv6 gateway-ip** to restore the default.

Syntax

```
apply ipv6 gateway-ip { ipv6-address | origin-nexthop }  
undo apply ipv6 gateway-ip
```

Default

The Gateway IP Address field is not set for EVPN IP prefix advertisement routes.

Views

Routing policy node view

Predefined user roles

network-admin

Parameters

ipv6-address: Specifies an IPv6 address.

origin-nexthop: Specifies the IPv6 address of the route next hop.

Usage guidelines

Use this command in an EVPN network to modify the Gateway IP Address field for EVPN IP prefix advertisement routes.

Examples

Set the Gateway IP Address field for EVPN IP prefix advertisement routes to the IPv6 address of the route next hop.

```
<Sysname> system-view
```

```
[Sysname] route-policy policy1 permit node 10
```

```
[Sysname-route-policy-policy1-10] apply ipv6 gateway-ip origin-nexthop
```

Related commands

```
apply gateway-ip
```

New feature: Enabling BGP to redistribute all IGP routes with the same destination address

Enabling BGP to redistribute all IGP routes with the same destination address

About this task

The ORIGIN attribute of IGP routes redistributed through this feature is INCOMPLETE.

Procedure (IPv4 unicast/IPv4 multicast)

1. Enter system view.
system-view
2. Enter BGP IPv4 unicast address family view or BGP-VPN IPv4 unicast address family view.
 - o Execute the following commands in sequence to enter BGP IPv4 unicast address family view:
bgp as-number [instance instance-name]
address-family ipv4 [unicast]

- Execute the following commands in sequence to enter BGP-VPN IPv4 unicast address family view:

```
bgp as-number [ instance instance-name ]
ip vpn-instance vpn-instance-name
address-family ipv4 [ unicast ]
```

3. Enable BGP to redistribute all IGP routes with the same destination address.

```
import-route multipath
```

By default, BGP does not redistribute IGP routes.

Procedure (IPv6 unicast/IPv6 multicast)

1. Enter system view.

```
system-view
```

2. Enter BGP IPv6 unicast address family view or BGP-VPN IPv6 unicast address family view.

- Execute the following commands in sequence to enter BGP IPv6 unicast address family view:

```
bgp as-number [ instance instance-name ]
address-family ipv6 [ unicast ]
```

- Execute the following commands in sequence to enter BGP-VPN IPv6 unicast address family view:

```
bgp as-number [ instance instance-name ]
ip vpn-instance vpn-instance-name
address-family ipv6 [ unicast ]
```

3. Enable BGP to redistribute all IGP routes with the same destination address.

```
import-route multipath
```

By default, BGP does not redistribute IGP routes.

Command reference

import-route multipath

Use **import-route multipath** to enable BGP to redistribute all IGP routes with the same destination address.

Use **undo import-route multipath** to restore the default.

Syntax

```
import-route multipath
```

```
undo import-route multipath
```

Default

BGP does not redistribute IGP routes.

Views

BGP IPv4 unicast address family view

BGP-VPN IPv4 unicast address family view

BGP IPv6 unicast address family view

BGP-VPN IPv6 unicast address family view

Predefined user roles

network-admin

Usage guidelines

The ORIGIN attribute of IGP routes redistributed through this command is INCOMPLETE.

Examples

In BGP IPv4 unicast address family view, enable BGP to redistribute all IGP routes with the same destination address.

```
<Sysname> system-view
```

```
[Sysname] bgp 100
```

```
[Sysname-bgp-default] address-family ipv4 unicast
```

```
[Sysname-bgp-default-ipv4] import-route multipath
```

New feature: Policing ND protocol packets

Policing ND protocol packets

The predefined control-plane QoS policy adds support for policing ND protocol packets. The default policing rate is 1000 pps.

Additionally, you can match the ND protocol packets and configure a traffic policing action for the matching ND protocol packets in a QoS policy. Then, you can apply the QoS policy to the inbound direction of the control plane to modify the default policing rate.

Command reference

Modified command: display qos policy control-plane pre-defined

Syntax

```
display qos policy control-plane pre-defined [ slot slot-number ]
```

Views

Any view

Change description

Before modification: The predefined control-plane QoS policy does not support ND protocol packet policing and no ND-related field appears in the command.

After modification: The predefined control-plane QoS policy adds support ND protocol packet policing and an ND-related field appears in the command.

Modified command: if-match

Syntax

```
if-match match-criteria
```

```
undo if-match match-criteria
```

Views

Traffic class view

Change description

Before modification: The `if-match control-plane protocol nd` command is not supported.

After modification: The `if-match control-plane protocol nd` command is supported.

New feature: Using an ACL to control TCP connections from IPv4 HTTP or HTTPS clients

Using an ACL to control TCP connections from IPv4 HTTP or HTTPS clients

As from this version, the device supports using an ACL to control TCP connections from IPv4 HTTP or HTTPS clients.

Command reference

http acl

Use `http acl` to control TCP connections from IPv4 HTTP clients.

Use `undo http acl` restore the default.

Syntax

```
http acl { advanced-acl-number | basic-acl-number }  
undo http acl
```

Default

TCP connections can be established from all IPv4 HTTP clients.

Views

System view

Predefined user roles

network-admin

Parameters

advanced-acl-number: Specifies an advanced IPv4 ACL number in the range of 3000 to 3999.

basic-acl-number: Specifies a basic IPv4 ACL number in the range of 2000 to 2999.

Usage guidelines

This command does not take effect on the existing TCP connections.

To control TCP connections from IPv4 HTTP clients, specify an existing ACL that has rules. Then, only IPv4 HTTP clients permitted by the ACL can establish TCP connections.

If no ACL is applied or the applied ACL does not exist or does not have rules, all IPv4 HTTP clients can establish TCP connections with the device.

Whether an IPv4 HTTP client can access a service after it establishes a TCP connection with the device depends on the configuration for this service. For example, for an IPv4 HTTP client to access the NETCONF over SOAP service, make sure the client is allowed to establish both TCP connections and NETCONF over SOAP connections with the device. You can use the `netconf soap acl` command to allow a client to establish a NETCONF over SOAP connection with the

device. For more information about the **netconf soap acl** command, see NETCONF in *Network Management and Monitoring Command Reference*.

Examples

Use ACL 2001 to allow only IPv4 HTTP clients from 1.1.1.1 to establish TCP connections with the device.

```
<Sysname> system-view
[Sysname] acl basic 2001
[Sysname-acl-ipv4-basic-2001] rule permit source 1.1.1.1 0
[Sysname-acl-ipv4-basic-2001] quit
[Sysname] http acl 2001
```

Related commands

netconf soap acl (*Network Management and Monitoring Command Reference*)

https acl

Use **https acl** to control TCP connections from IPv4 HTTPS clients.

Use **undo https acl** restore the default.

Syntax

```
https acl { advanced-acl-number | basic-acl-number }
undo https acl
```

Default

TCP connections can be established from all IPv4 HTTPS clients.

Views

System view

Predefined user roles

network-admin

Parameters

advanced-acl-number: Specifies an advanced IPv4 ACL number in the range of 3000 to 3999.

basic-acl-number: Specifies a basic IPv4 ACL number in the range of 2000 to 2999.

Usage guidelines

This command does not take effect on the existing TCP connections.

To control TCP connections from IPv4 HTTPS clients, specify an existing ACL that has rules. Then, only IPv4 HTTPS clients permitted by the ACL can establish TCP connections with the device.

If no ACL is applied or the applied ACL does not exist or does not have rules, all IPv4 clients can establish TCP connections.

Whether an IPv4 HTTPS client can access a service after it establishes a TCP connection with the device depends on the configuration for this service. For example, for an IPv4 HTTPS client to access the NETCONF over SOAP service, make sure the client is allowed to establish both TCP connections and NETCONF over SOAP connections with the device. You can use the **netconf soap acl** command to allow a client to establish a NETCONF over SOAP connection with the device. For more information about the **netconf soap acl** command, see NETCONF in *Network Management and Monitoring Command Reference*.

Examples

Use ACL 2001 to allow only IPv4 HTTPS clients from 1.1.1.1 to establish TCP connections with the device.

```
<Sysname> system-view
[Sysname] acl basic 2001
[Sysname-acl-ipv4-basic-2001] rule permit source 1.1.1.1 0
[Sysname-acl-ipv4-basic-2001] quit
[Sysname] https acl 2001
```

Related commands

netconf soap acl (*Network Management and Monitoring Command Reference*)

New feature: Enabling Layer 3 packet statistics counting

Enabling Layer 3 packet statistics counting

About this task

The Layer 3 packet statistics counting feature counts statistics about incoming and outgoing IP packets on an interface. To display the collected statistics, execute the **display ipv6 statistics** or **display ip statistics** command.

Restrictions and guidelines

When the interface is processing a large number of packets, the Layer 3 packet statistics counting will cause high CPU usage and degrade the forwarding performance. If the statistics are not necessary, disable this feature to ensure the device performance.

Procedure

1. Enter system view.
system-view
2. Enter interface view.
interface *interface-type interface-number*
3. Enable Layer 3 packet statistics counting.
statistics l3-packet enable { inbound | outbound }
By default, Layer 3 packet statistics counting is disabled.

Command reference

statistics l3-packet enable

Use **statistics l3-packet enable** to enable Layer 3 packet statistics counting.

Use **undo statistics l3-packet enable** to disable Layer 3 packet statistics counting.

Syntax

```
statistics l3-packet enable { inbound | outbound }
undo statistics l3-packet enable { inbound | outbound }
```

Default

Layer 3 packet statistics counting is disabled.

Views

Interface view

Predefined user roles

network-admin

Parameters

inbound: Enables statistics counting for incoming Layer 3 packets.

outbound: Enables statistics counting for outgoing Layer 3 packets.

Usage guidelines

The Layer 3 packet statistics counting feature counts statistics of incoming and outgoing IP packets on an interface. To display the collected statistics, execute the **display ipv6 statistics** or **display ip statistics** command. When the interface is processing a large number of packets, enabling this feature will cause high CPU usage and degrade the forwarding performance. If the statistics are not necessary, disable this feature to ensure the device performance.

Examples

Enable Layer 3 packet statistics counting on Ten-GigabitEthernet 1/0/1.

```
<Sysname> system-view
[Sysname] interface ten-gigabitethernet 1/0/1
[Sysname-Ten-GigabitEthernet 1/0/1] statistics l3-packet enable inbound
[Sysname-Ten-GigabitEthernet 1/0/1] statistics l3-packet enable outbound
```

Related commands

display ip interface

display ip statistics

display ipv6 statistics

New feature: Specifying the PKI domain for secure communication with collectors

Specifying the PKI domain for secure communication with collectors

About this task

By default, the gRPC connection between the device and a collector does not provide data encryption service or require authentication. After you specify a PKI domain, the device and the collector will use TLS for data encryption and bidirectional certificate-based authentication to improve communication security.

Restrictions and guidelines

The specified PKI domain must already exist and has correct certificate and key settings. For more information about PKI domains, see PKI configuration in *Security Configuration Guide*.

After you specify the PKI domain, the gRPC service will reboot, closing the connections to collectors. The collectors must re-initiate the connections.

Procedure

1. Enter system view.
system-view
2. Specify the PKI domain for secure communication with collectors.

grpc pki domain *domain-name*

By default, no PKI domain is specified for secure communication with collectors.

Command reference

grpc pki domain

Use **grpc pki domain** to specify the PKI domain for secure gRPC connections to collectors.

Use **undo grpc pki domain** to restore the default.

Syntax

grpc pki domain *domain-name*

undo grpc pki domain

Default

No PKI domain is specified for secure gRPC connections to collectors.

Views

System view

Predefined user roles

network-admin

Parameters

domain-name: Specifies a PKI domain by its name, a case-insensitive string of 1 to 31 characters.

Usage guidelines

By default, the gRPC connection between the device and a collector does not provide data encryption service or require authentication. After you specify a PKI domain, the device and the collector will use TLS for data encryption and bidirectional certificate-based authentication to improve communication security.

For the device to establish gRPC connections to collectors successfully, make sure the PKI domain already exists and has correct certificate and key settings. For more information about PKI configuration, see *Security Configuration Guide*.

Executing this command reboots the gRPC service and closes all gRPC connections to collectors. In dial-in mode, the collectors must re-initiate the connections. In dial-out mode, the device will automatically re-initiate the connections.

Examples

Specify the PKI domain for secure gRPC connections to collectors.

```
<Sysname> system-view
```

```
[Sysname] grpc pki domain grpc_test
```

New feature: DRNI MAD

Configuring DRNI MAD

About this task

DRNI MAD configuration methods

When you configure DRNI MAD, use either of the following methods:

- To shut down all network interfaces on the secondary DR member device except a few special-purpose interfaces that must be retained in up state:
 - Set the default DRNI MAD action to DRNI MAD DOWN. For more information, see ["Configuring the default DRNI MAD action on network interfaces."](#)
 - Exclude interfaces from being shut down by DRNI MAD.

This method is applicable to most network environments.
- To have the secondary DR member device retain a large number of interfaces in up state and shut down the remaining interfaces:
 - Set the default DRNI MAD action to NONE. For more information, see ["Configuring the default DRNI MAD action on network interfaces."](#)
 - Specify network interfaces that must be shut down by DRNI MAD. For more information, see ["Specifying interfaces to be shut down by DRNI MAD when the DR system splits."](#)

One applicable scenario of this method is the EVPN environment in which you use a VXLAN tunnel as the IPL. In this scenario, you must retain a large number of logical interfaces (for example, tunnel and loopback interfaces) in up state.

List of automatically included interfaces

DRNI MAD will always shut down the ports in the system-configured included port list if the device acts as the secondary DR member device when the DR system splits.

This list contains aggregation member ports of DR interfaces. To identify system-configured included ports, execute the `display drni mad verbose` command.

List of automatically excluded interfaces

DRNI MAD will not shut down the ports in the following list when the DR system splits:

- System-configured excluded port list in DRNI MAD:
 - IPP.
 - Aggregation member interfaces if a Layer 2 aggregate interface is used as the IPP.
 - DR interfaces.
 - Management interfaces.

To identify these interfaces, execute the `display drni mad verbose` command.
- Interfaces manually or automatically excluded from being shut down by IRF MAD. To identify these interfaces, execute the `display mad verbose` command.
- Network interfaces used for special purposes, including:
 - Interfaces placed in a loopback test by using the `loopback` command.
 - Interfaces in a service loopback group.
 - Interfaces in a mirroring group.
 - Interfaces forced to stay up by using the `port-up mode` command.

Configuring the default DRNI MAD action on network interfaces

About this task

You can configure DRNI MAD to take either of the following default actions on network interfaces if the device acts as the secondary DR member device when the DR system splits:

- **DRNI MAD DOWN**—DRNI MAD will shut down all network interfaces on the secondary DR member device when the DR system splits, except the interfaces excluded manually or by the system.
- **NONE**—DRNI MAD will not shut down any network interfaces when the DR system splits, except the interfaces configured manually or by the system to be shut down by DRNI MAD.

Restrictions and guidelines

The DRNI MAD DOWN action will not take effect on the interfaces listed in "[List of automatically excluded interfaces](#)."

The DRNI MAD DOWN action will always take on the interfaces listed in "[List of automatically included interfaces](#)," even if the default DRNI MAD action is NONE.

Procedure

1. Enter system view.
system-view
2. Configure the default DRNI MAD action to take on network interfaces on the secondary DR member device when the DR system splits.
drni mad default-action { down | none }
By default, DRNI MAD shuts down network interfaces on the secondary DR member device.

Specifying interfaces to be shut down by DRNI MAD when the DR system splits

About this task

By default, DRNI MAD automatically shuts down the interfaces listed in "[List of automatically included interfaces](#)" if the device is the secondary DR member device when the DR system splits.

To specify additional interfaces to be shut down by DRNI MAD, perform this task.

You typically perform this task when the default DRNI MAD action is set to NONE.

Restrictions and guidelines

The DRNI MAD DOWN action will not take effect on the interfaces listed in "[List of automatically excluded interfaces](#)."

Procedure

1. Enter system view.
system-view
2. Specify interfaces to be shut down by DRNI MAD when the DR system splits.
drni mad include interface *interface-type interface-number*
By default, the user-configured included port list does not contain any ports.

Command reference

New command: drni mad default-action

Use **drni mad default-action** to configure the default DRNI MAD action to take on network interfaces on the secondary DR member device when the DR system splits.

Use **undo drni mad default-action** to restore the default.

Syntax

```
drni mad default-action { down | none }  
undo drni mad default-action
```

Default

DRNI MAD shuts down network interfaces on the secondary DR member device.

Views

System view

Predefined user roles

network-admin

Parameters

down: Sets the action to DRNI MAD DOWN. DRNI MAD will shut down all network interfaces on the secondary DR member device when the DR system splits, except the interfaces excluded manually or by the system.

none: Sets the DRNI MAD action to NONE. DRNI MAD will not shut down any network interfaces when the DR system splits, except the interfaces configured manually or by the system to be shut down by DRNI MAD.

Usage guidelines

In most network environments, use the DRNI MAD DOWN action in conjunction with the **drni mad exclude interface** command.

- The DRNI MAD DOWN action helps avoid network issues caused by multi-active collision, which occurs if the IPL goes down while the keepalive link is up.
- The **drni mad exclude interface** command enables you to exclude special-purpose interfaces, for example, interfaces used in DRNI for DR keepalive detection.

This configuration method is inefficient if a large number of interfaces on the secondary DR member must be retained in up state after the DR system splits. For example, if you use a VXLAN tunnel as the IPL in an EVPN environment, you must retain a large number of logical interfaces (for example, tunnel and loopback interfaces) in up state.

In these situations, use the following method to configure DRNI MAD:

- Set the default DRNI MAD action to NONE.
- Execute the **drni mad include interface** command to specify interfaces that must be shut down by DRNI MAD in addition to those already automatically specified by the system.

The DRNI MAD DOWN action will not take effect on the following interfaces:

- Interfaces automatically or manually excluded from being shut down by DRNI MAD. To identify these interfaces, execute the **display drni mad verbose** command.
- Interfaces automatically or manually excluded from being shut down by IRF MAD. To identify these interfaces, execute the **display mad verbose** command.
- Interfaces used for special purposes, including:
 - Interfaces placed in a loopback test by using the **loopback** command.
 - Interfaces in a service loopback group.
 - Interfaces in a mirroring group.
 - Interfaces forced to stay up by using the **port-up mode** command.

Examples

Configure DRNI MAD to shut down all network interfaces except excluded interfaces when the DR system splits.

```
<Sysname> system-view
```

```
[Sysname] drni mad default-action down
```

Related commands

display drni mad verbose

display mad (*Virtual Technologies Command Reference*)

```
drni mad exclude interface
drni mad include interface
loopback (Interface Command Reference)
port up-mode (Interface Command Reference)
```

New command: drni mad include interface

Use **drni mad include interface** to add an interface to the user-configured list of included ports, which will be shut down by DRNI MAD when the DR system splits.

Use **undo drni mad include interface** to remove a network interface from the user-configured list of included ports.

Syntax

```
drni mad include interface interface-type interface-number
undo drni mad include interface interface-type interface-number
```

Default

The user-configured included port list does not contain any ports.

Views

System view

Predefined user roles

network-admin

Parameters

interface-type interface-number: Specifies an interface by its type and number.

Usage guidelines

This command is typically used when the default DRNI MAD action is set to NONE. In this situation, DRNI MAD does not shut down the network interfaces on the secondary DR member device when the DR system splits. To avoid network issues caused by multi-active collision, add network interfaces to the user-configured include port list if they are not in the system-configured included port list. DRNI MAD on the secondary DR member device will shut down the interfaces in the system-configured and user-configured included port lists when the DR system splits.

For example, if you use a VXLAN tunnel as the IPL in an EVPN environment, you must retain a large number of logical interfaces (for example, tunnel and loopback interfaces) in up state. To improve configuration efficiency, you can set the default DRNI MAD action to NONE and add interfaces to the user-configured include port list.

The DRNI MAD DOWN action will not take effect on the following interfaces:

- Network interfaces automatically or manually excluded from being shut down by DRNI MAD. To identify these interfaces, execute the **display drni mad verbose** command.
- Network interfaces automatically or manually excluded from being shut down by IRF MAD. To identify these interfaces, execute the **display mad verbose** command.
- Network interfaces used for special purposes, including:
 - Interfaces placed in a loopback test by using the **loopback** command.
 - Interfaces in a service loopback group.
 - Interfaces in a mirroring group.
 - Interfaces forced to stay up by using the **port-up mode** command.

Examples

```
# Configure DRNI MAD to shut down Bridge-Aggregation 1 on the secondary DR member device
when the DR system splits.
```

```
<Sysname> system-view
```

```
[Sysname] drni mad include interface bridge-aggregation 1
```

Related commands

```
display drni mad verbose
```

```
display mad (Virtual Technologies Command Reference)
```

```
drni mad default-action
```

```
loopback (Interface Command Reference)
```

```
mirroring-group reflector-port (Network Management and Monitoring Command Reference)
```

```
port service-loopback group
```

```
port up-mode (Interface Command Reference)
```

Modified command: display drni mad verbose

Syntax

```
display drni mad verbose
```

Views

Any view

Change description

Before modification: This command does not display the interfaces placed in DRNI MAD DOWN state after the DR system splits.

```
# Display detailed DRNI MAD information.
```

```
<Sysname> display drni mad verbose
```

```
DRNI MAD DOWN state      : No
```

```
Restore delay            : 30 s
```

```
Keepalive status         : Normal
```

```
System number  Keepalive Destination IP      UDP port
```

```
1              1.1.1.1                       1000
```

```
Excluded ports(user-configured):
```

```
Excluded ports(system-configured):
```

```
Management interfaces:
```

```
  M-GigabitEthernet0/0/0
```

```
DR interfaces:
```

```
  Bridge-Aggregation4
```

```
IPP:
```

```
  Bridge-Aggregation3
```

```
Member interfaces of IPP Bridge-Aggregation3:
```

```
  Ten-GigabitEthernet1/0/1
```

```
  Ten-GigabitEthernet1/0/2
```

After modification: This command displays the interfaces placed in DRNI MAD DOWN state after the DR system splits.

```
# Display detailed DRNI MAD information.
```

```

<Sysname> display drni mad verbose
DRNI MAD DOWN state      : No
Restore delay            : 30 s
Keepalive status         : Normal
System number   Keepalive Destination IP      UDP port
1                1.1.1.1                      1000
DRNI MAD default action: DOWN
Port configuration for DRNI MAD DOWN action:
Included ports(user-configured):
  Bridge-Aggregation1
  Ten-GigabitEthernet1/0/5(inactive)
Included ports(system-configured):
  Member interfaces of DR Bridge-Aggregation3:
    Ten-GigabitEthernet1/0/3
Excluded ports(user-configured):
  Bridge-Aggregation2
  Ten-GigabitEthernet1/0/4
Excluded ports(system-configured):
  Management interfaces:
    M-GigabitEthernet0/0/0
  DR interfaces:
    Bridge-Aggregation4
  IPP:
    Bridge-Aggregation3
  Member interfaces of IPP Bridge-Aggregation3:
    Ten-GigabitEthernet1/0/1
    Ten-GigabitEthernet1/0/2

```

Table 2 Command output

| Field | Description |
|-----------------------------------|---|
| DRNI MAD default action | <p>Default action to take on network interfaces when the DR system splits:</p> <ul style="list-style-type: none"> DOWN—Shut down interfaces and place them in DRNI MAD DOWN state. NONE—DRNI MAD does not take action on interfaces. |
| Included ports(user-configured) | <p>Network interfaces manually configured to be shut down by DRNI MAD when the DR system splits.</p> <p>An interface entry will be marked as ineffective if the system does not allow the specified interface to be shut down by DRNI MAD. DRNI MAD will not shut down interfaces in ineffective entries when the DR system splits.</p> <p>The following are interfaces not allowed to be shut down by DRNI MAD:</p> <ul style="list-style-type: none"> Interfaces automatically excluded from being shut down by DRNI MAD. Interfaces manually or automatically excluded from being shut down by IRF MAD. Interfaces used for special purposes. |
| Included ports(system-configured) | <p>Network interfaces automatically set by the system to shut down by DRNI MAD when the DR system splits.</p> <p>Aggregation member ports of DR interfaces are in this category of interfaces.</p> |

New feature: Configuring support for forwarding DHCPv6 replies based on MAC address table

Configuring support for forwarding DHCPv6 replies based on MAC address table

About this task

In a distributed EVPN gateway network, VSI interfaces of all distributed EVPN gateways have the same IP address, and the DHCPv6 relay agent is enabled on EVPN gateways. When a DHCPv6 client sends a request to its connected EVPN gateway, the gateway records the request forwarding information for the client before relaying the request to the DHCPv6 server. The request forwarding information contains the MAC address of the DHCPv6 client and output interface of the request packet. If another EVPN gateway receives the reply for this request, the gateway discards the reply by default because it does not have the matching forwarding information for the reply. As a result, the client cannot obtain an IPv6 address or prefix.

To solve this problem, configure the support for forwarding replies based on MAC address table on the relay agent that is connected to the DHCPv6 server. This feature allows the relay agent to look up the MAC address table for the output interface for the replies that do not have the request forwarding information. This feature ensures that DHCPv6 clients receive DHCPv6 reply packets.

If the **broadcast** keyword is specified, the DHCPv6 relay agent broadcasts the DHCPv6 reply out of the interface found in the matching MAC address entry. When the DHCPv6 relay agent connecting to the requesting DHCPv6 client receives the reply, it sends the reply to its CPU before broadcasting the reply. If recording relay entries is enabled on this relay agent, a relay entry is generated for the DHCPv6 client.

If the **broadcast** keyword is not specified, the DHCPv6 relay agent connecting to the requesting client will directly forward the reply based on the MAC address instead of delivering it up to the CPU. In this case, the DHCPv6 relay agent cannot be aware that the packet is a DHCPv6 reply, thus no relay entry is generated for the client even if the recording of relay entries is enabled.

Procedure

1. Enter system view.

```
system-view
```

2. Enable MAC address table lookup for DHCPv6 replies that do not have request forwarding information.

```
ipv6 dhcp relay mac-forward enable [ broadcast ]
```

By default, the DHCPv6 relay agent discards a DHCPv6 reply if the relay agent does not have the request forwarding information for the reply.

Command reference

ipv6 dhcp relay mac-forward enable

Use **ipv6 dhcp relay mac-forward enable** to enable MAC address table lookup for DHCPv6 replies that do not have forwarding information on the relay agent.

Use **undo ipv6 dhcp relay mac-forward enable** to restore the default.

Syntax

```
ipv6 dhcp relay mac-forward enable [ broadcast ]
```

```
undo ipv6 dhcp relay mac-forward enable
```

Default

The DHCPv6 relay agent discards a DHCPv6 reply if the agent does not have the forwarding information for the reply.

Views

System view

Predefined user roles

network-admin

Parameters

broadcast: Broadcasts DHCPv6 replies out of the interface found in the matching MAC address entry. If you do not specify this keyword, the DHCPv6 relay agent unicasts DHCPv6 replies.

Usage guidelines

If you execute this command multiple times, the most recent configuration takes effect.

Examples

Enable MAC address table lookup for DHCPv6 replies that do not have forwarding information on the relay agent.

```
<Sysname> system-view
```

```
[Sysname] ipv6 dhcp relay mac-forward enable
```

New feature: Configuring basic DHCPv6 snooping features on a VXLAN network

Configuring basic DHCPv6 snooping features on a VXLAN network

About this task

On a VXLAN network, you can configure the following interfaces as DHCPv6 snooping trusted interfaces:

- ACs that are mapped to a VSI.
- VXLAN tunnel interfaces that are assigned to a VSI.

When the VTEP with DHCPv6 snooping configured receives a DHCPv6 request, the VTEP forwards this request through the trusted ACs or VXLAN tunnel interfaces. Untrusted ACs or VXLAN tunnel interface on the VTEP drop incoming DHCPv6 replies.

Restrictions and guidelines

If the DHCPv6 server is in the local site, configure the AC that connects to the DHCPv6 server as trusted. If the DHCPv6 server is in a remote site, configure the VXLAN tunnel interface as trusted.

Configuring basic DHCPv6 snooping features in a VXLAN network (DHCPv6 server at the local site)

1. Enter system view.
system-view
2. Enable DHCPv6 snooping globally.
ipv6 dhcp snooping enable
By default, DHCPv6 snooping is disabled globally.
3. Enter interface view.

interface *interface-type interface-number*

4. Enter Ethernet service instance view.

service-instance *instance-id*

5. Configure the AC as the DHCPv6 snooping trusted interface.

ipv6 dhcp snooping trust

By default, all ports are untrusted after DHCPv6 snooping is enabled.

6. (Optional.) Enable recording of client information in DHCPv6 snooping entries on the ACs mapped to the VSI and VXLAN tunnel interfaces assigned to the VSI.

- a. Exit to interface view.

quit

- b. Exit to the system view.

quit

- c. Enter VSI view.

vsi *vsi-name*

- d. Enable recording of client information in DHCP snooping entries on the ACs mapped to the VSI and VXLAN tunnel interfaces assigned to the VSI.

ipv6 dhcp snooping binding record

By default, the recording of DHCP snooping entries is disabled.

Configuring basic DHCPv6 snooping features in a VXLAN network (DHCPv6 server at a remote site)

1. Enter system view.

system-view

2. Enable DHCPv6 snooping globally.

ipv6 dhcp snooping enable

By default, DHCP snooping is disabled globally.

3. Enter VSI view.

vsi *vsi-name*

4. Configure the VXLAN tunnel interfaces as the DHCP snooping trusted interface.

ipv6 dhcp snooping trust tunnel

By default, all ports are untrusted after DHCPv6 snooping is enabled.

This command sets all VXLAN tunnel interfaces in the VSI as DHCPv6 snooping trusted interfaces.

5. (Optional.) Enable recording of client information in DHCP snooping entries on the ACs mapped to the VSI and VXLAN tunnel interfaces assigned to the VSI.

ipv6 dhcp snooping binding record

By default, the recording of DHCP snooping entries is disabled.

Command changes

ipv6 dhcp snooping binding record

Use **ipv6 dhcp snooping binding record** to enable recording DHCPv6 snooping address entries.

Use **undo ipv6 dhcp snooping binding record** to disable recording DHCPv6 snooping address entries.

Syntax

```
ipv6 dhcp snooping binding record
undo ipv6 dhcp snooping binding record
```

Default

Recording of DHCPv6 snooping address entries is disabled.

Views

VSI view

Predefined user roles

network-admin

Usage guidelines

If you configure this command in a VSI view, it takes effect on the ACs that are mapped to the VSI and the VXLAN tunnel interfaces that are assigned to the VSI.

Examples

```
# Enable recording DHCPv6 snooping address entries for VSI a.
<Sysname> system-view
[Sysname] vsi a
[Sysname-vsi-a] ipv6 dhcp snooping binding record
```

ipv6 dhcp snooping trust tunnel

Use **ipv6 dhcp snooping trust tunnel** to configure VXLAN tunnel interfaces assigned to a VSI as trusted interfaces.

Use **undo ipv6 dhcp snooping trust tunnel** to restore the default.

Syntax

```
ipv6 dhcp snooping trust tunnel
undo ipv6 dhcp snooping trust tunnel
```

Default

After you enable DHCPv6 snooping, all tunnel interfaces are untrusted.

Views

VSI view

Predefined user roles

network-admin

Usage guidelines

With DHCPv6 snooping enabled, only the DHCPv6 snooping trusted ports forward responses from DHCPv6 servers. To ensure VXLAN tunnel interfaces can forward DHCPv6 messages correctly, configure the VXLAN tunnel interfaces as DHCPv6 snooping trusted.

Examples

```
# Configure the tunnel interfaces as trusted in VSI a.
<Sysname> system-view
[Sysname] vsi a
[Sysname-vsi-a] ipv6 dhcp snooping trust tunnel
```

Related commands

```
display ipv6 dhcp snooping trust
```

Modified command: ipv6 dhcp snooping trust

Syntax

```
ipv6 dhcp snooping trust
undo ipv6 dhcp snooping trust
```

Views

Ethernet service instance view

Change description

Before modification: This command is not available in Ethernet service instance view.

After modification: This command is available in Ethernet service instance view.

New feature: Configuring BGP VPNv6 routes to use private network next hops

Configuring BGP VPNv6 routes to use private network next hops

About this task

By default, the device does not change the next hop attribute of a received BGP VPNv6 route. The next hop address of a BGP VPNv6 route is a public address. This feature changes the next hop of a BGP VPNv6 route received from a peer or peer group to an IP address in the VPN instance. The outgoing label of the VPNv6 route is also changed to an invalid value. For example, the device received a VPNv6 route and its next hop address is 10.1.1.1, which is a public address by default. After this feature is configured, the next hop address changes to private address 10.1.1.1.

Restrictions and guidelines

After you configure this feature, the following applies:

- The device re-establishes the BGP sessions to the specified peer or to all peers in the specified peer group.
- The device receives a BGP VPNv6 route only when its RD is the same as a local RD.
- When advertising a BGP VPNv6 route received from the specified peer or peer group, the device does not change the route target attribute of the route.
- If you delete a VPN instance or its RD, BGP VPNv6 routes received from the specified peer or peer group and in the VPN instance will be deleted.

Procedure

1. Enter system view.
system-view
2. Enter BGP instance view.
bgp *as-number* [**instance** *instance-name*]
3. Enter BGP VPNv6 address family view.
address-family **vpnvp6**
4. Change the next hop of a BGP VPNv6 route received from a peer or peer group to a VPN instance address.

peer { *group-name* | *ipv4-address* [*mask-length*] } **next-hop-vpn**

By default, the device does not change the next hop attribute of a received BGP VPNv6 route, and the next hop belongs to the public network.

Command reference

peer next-hop-vpn

Use **peer next-hop-vpn** to change the next hop of a BGP VPNv6 route received from a peer or peer group to a VPN instance address.

Use **undo peer next-hop-vpn** to restore the default.

Syntax

peer { *group-name* | *ipv4-address* [*mask-length*] } **next-hop-vpn**

undo peer { *group-name* | *ipv4-address* [*mask-length*] } **next-hop-vpn**

Default

The device does not change the next hop attribute of a received BGP VPNv6 route, and the next hop belongs to the public network.

Views

BGP VPNv6 address family view

Predefined user roles

network-admin

Parameters

group-name: Specifies a peer group by its name, a case-sensitive string of 1 to 47 characters. The specified group must have been created.

ipv4-address: Specifies a peer by its IP address. The specified peer must have been created.

mask-length: Specifies a mask length in the range of 0 to 32. You can use the *ipv4-address* and *mask-length* arguments together to specify a subnet. If you specify a subnet, this command changes the next hop of BGP VPNv6 routes received from the dynamic peers in the subnet.

Usage guidelines

By default, the device does not change the next hop attribute of a received BGP VPNv6 route. The next hop address of a BGP VPNv6 route is a public address. This command changes the next hop address of a BGP VPNv6 route received from a peer or peer group to a VPN instance address. The outgoing label of the VPNv6 route is also changed to an invalid value. For example, the device received a VPNv6 route and its next hop address is 10.1.1.1, which is a public address by default. After this command is executed, the next hop address changes to private address 10.1.1.1.

After this command is executed, the following applies:

- The device re-establishes the BGP sessions to the specified peer or to all peers in the specified peer group.
- The device receives a BGP VPNv6 route only when its RD is the same as a local RD.
- When advertising a BGP VPNv6 route received from the specified peer or peer group, the device does not change the route target attribute of the route.
- If you delete a VPN instance or its RD, BGP VPNv6 routes received from the specified peer or peer group and in the VPN instance will be deleted.

Examples

In BGP VPNv6 address family view, change the next hop of BGP VPNv6 routes received from peer 1.1.1.1 to a VPN instance address.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] address-family vpnv6
[Sysname-bgp-default-vpnv6] peer 1.1.1.1 next-hop-vpn
```

New feature: Configuring BGP route replication between public and VPN instances

Configuring BGP route replication between public and VPN instances

About this task

In a BGP/MPLS L3VPN network, only VPN instances that have matching route targets can communicate with each other.

In traffic cleaning scenarios, traffic between the public and private networks are filtered by firewalls and traffic of different tenants is assigned to different VPNs. To enable the tenants to communicate with the public network under the protection of firewalls, you can configure BGP route replication between public and VPN instances.

After you configure this feature, the public and VPN instances replicate all BGP routes including route attributes from each other.

Restrictions and guidelines

This feature also enables BGP route replication between VPN instances, so VPNs cannot be isolated. Configure this feature only in specific scenarios, for example, the traffic cleaning scenario.

Procedure

1. Enter system view.
system-view
2. Enter BGP instance view.
bgp as-number [instance instance-name]
3. Enable BGP route replication between public and VPN instances.
route-replicate enable
By default, BGP route replication between public and VPN instances is disabled.
4. Return to system view.
quit
5. Enter public instance view, public instance IPv4 address family view, or public instance IPv6 address family view.
 - o Enter public instance view.
ip public-instance
 - o Execute the following commands in sequence to enter public instance IPv4 address family view:
ip public-instance

address-family ipv4

- Execute the following commands in sequence to enter public instance IPv6 address family view:

ip public-instance

address-family ipv6

6. Apply an import routing policy to the public instance.

import route-policy *route-policy*

By default, all routes matching the import target attribute are accepted.

7. Apply an export routing policy to the public instance.

export route-policy *route-policy*

By default, routes to be advertised are not filtered.

Command reference

route-replicate enable

Use **route-replicate enable** to enable BGP route replication between public and VPN instances.

Use **undo route-replicate enable** to disable BGP route replication between public and VPN instances.

Syntax

route-replicate enable

undo route-replicate enable

Default

BGP route replication between public and VPN instances is disabled.

Views

BGP instance view

Predefined user roles

network-admin

Usage guidelines

In a BGP/MPLS L3VPN network, only VPN instances that have matching route targets can communicate with each other.

After you execute this command on PEs, the public and VPN instance replicate all BGP routes including the BGP route attributes from each other. As a result, the public and VPN instances can communicate with other.

This command also enables BGP route replication between VPN instances, so VPNs cannot be isolated. Use this command only in specific scenarios, for example, the traffic cleaning scenario.

In traffic cleaning scenarios, traffic between the public and private networks are filtered by firewalls and traffic of different tenants is assigned to different VPNs. To enable the tenants to communicate with the public network under the protection of firewalls, you can configure BGP route replication between public and VPN instances.

Examples

In BGP instance view, enable mutual route replication between public and VPN instances.

```
<Sysname> system-view
```

```
[Sysname] bgp 100
```

```
[Sysname-bgp-default] route-replicate enable
```

export route-policy

Use **export route-policy** to apply an export routing policy to the public instance or a VPN instance.

Use **undo export route-policy** to restore the default.

Syntax

```
export route-policy route-policy
```

```
undo export route-policy
```

Default

No export routing policy is applied to the public instance or a VPN instance.

Views

Public instance view

Public instance IPv4 address family view

Public instance IPv6 address family view

Predefined user roles

network-admin

Parameters

route-policy: Specifies a routing policy by its name, a case-sensitive string of 1 to 63 characters.

Usage guidelines

You can specify an export routing policy to filter advertised routes or modify their route attributes for the public instance or the VPN instance.

If you execute this command multiple times, the most recent configuration takes effect.

An export routing policy specified in VPN instance view or public instance view applies to all address families in the VPN instance or public instance. An export routing policy specified in an address family view applies only to the address family.

An address family prefers the export routing policy specified in the address family view over the one specified in VPN instance view or public instance view.

Examples

Apply export routing policy **poly-1** to the public instance.

```
<Sysname> system-view
```

```
[Sysname] ip public-instance
```

```
[Sysname-public-instance] export route-policy poly-1
```

Related commands

```
import route-policy
```

```
route-policy (Layer 3—IP Routing Command Reference)
```

import route-policy

Use **import route-policy** to apply an import routing policy to the public instance or a VPN instance.

Use `undo import route-policy` to restore the default.

Syntax

```
import route-policy route-policy  
undo import route-policy
```

Default

All routes matching the import target attribute are accepted.

Views

Public instance view

Public instance IPv4 address family view

Public instance IPv6 address family view

Predefined user roles

network-admin

Parameters

route-policy: Specifies a routing policy by its name, a case-sensitive string of 1 to 63 characters.

Usage guidelines

You can specify an import routing policy to filter received routes or modify their route attributes for the public instance or a VPN instance.

If you execute this command multiple times, the most recent configuration takes effect.

An import routing policy specified in VPN instance view or public instance view applies to all address families in the VPN instance or public instance. An import routing policy specified in an address family view applies only to the address family.

An address family prefers the import routing policy specified in the address family view over the one specified in VPN instance view or public instance view.

Examples

```
# Apply import routing policy poly-1 to the public instance.  
<Sysname> system-view  
[Sysname] ip public-instance  
[Sysname-public-instance] import route-policy poly-1
```

Related commands

```
export route-policy  
route-policy (Layer 3—IP Routing Command Reference)
```

New feature: Displaying IPv6 multicast group information that MLD snooping learns from EVPN

Displaying IPv6 multicast group information that MLD snooping learns from EVPN

Execute the `display` command in any view.

| Task | Command |
|--|---|
| Display IPv6 multicast group information that MLD snooping learns from EVPN. | display mld-snooping evpn-group [<i>ipv6-group-address</i> <i>ipv6-source-address</i>] * [vsi <i>vsi-name</i>] [slot <i>slot-number</i>] |

Command reference

display mld-snooping evpn-group

Use **display mld-snooping evpn-group** to display IPv6 multicast group information that MLD snooping learns from EVPN.

Syntax

```
display      mld-snooping      evpn-group      [      ipv6-group-address      |
ipv6-source-address ] * [ vsi vsi-name ] [ slot slot-number ]
```

Views

Any view

Predefined user roles

network-admin
network-operator

Parameters

ipv6-group-address: Specifies an IPv6 multicast group address. The value range for this argument is FFxy::/16 (excluding FFx1::/16 and FFx2::/16), where "x" and "y" represent any hexadecimal numbers in the range of 0 to F. If you do not specify an IPv6 multicast group, this command displays information that MLD snooping learns from EVPN for all IPv6 multicast groups.

ipv6-source-address: Specifies an IPv6 multicast source address. If you do not specify an IPv6 multicast source, this command displays information that MLD snooping learns from EVPN for all IPv6 multicast sources.

vsi *vsi-name*: Specifies a VSI by its name, a case-sensitive string of 1 to 31 characters. If you do not specify a VSI, this command displays multicast group information that MLD snooping learns from EVPN for all VSIs.

slot *slot-number*: Specifies an IRF member device by its member ID. If you do not specify a card, this command displays multicast group information that MLD snooping learns from EVPN on the master device.

New feature: Enabling password control blacklist for all user line types

Enabling password control blacklist for all user line types

About this task

After you enable the global password control feature, the password control blacklist feature, by default, is enabled only for FTP and VTY users. For the password control blacklist feature for login users on all user lines (AUX and VTY lines) to take effect, you must also enable this feature.

Procedure

1. Enter system view.
system-view
2. Enable the global password control feature.
In non-FIPS mode:
password-control enable [network-class]
By default, the global password control feature is disabled for device management and network access users.
In FIPS mode:
password-control enable [network-class]
By default, the global password control feature is enabled for device management users and cannot be disabled. The global password control feature is disabled for network access users.
3. Enable the password control blacklist feature for all user line types.
password-control blacklist all-line
By default, the password control blacklist feature is disabled for all user line types. The password control blacklist feature is enabled only for FTP and VTY users after you enable the global password control feature.

Command changes

password-control blacklist all-line

Use **password-control blacklist all-line** to enable the password control blacklist feature for all user line types.

Use **undo password-control blacklist all-line** to disable the password control blacklist feature for all user line types.

Syntax

password-control blacklist all-line
undo password-control blacklist all-line

Default

The password control blacklist feature is disabled for all user line types.

Views

System view

Predefined user roles

network-admin

Usage guidelines

By default, the password control blacklist feature is enabled only for FTP and VTY users when you enable the global password control feature. To enable the password control blacklist feature for login users on all user lines (AUX and VTY lines), you must also execute this command. For more information about user lines, see login management configuration in *Fundamentals Configuration Guide*.

With this feature enabled, the number of consecutive login failures on the same user account is accumulated for users on AUX lines. If the number reaches the limit for AUX users, the system adds the user account to the blacklist and takes corresponding action.

Use **password-control login-attempt** to configure the maximum number of consecutive login failures and the action to be taken when the maximum number is reached.

Examples

Enable the password control blacklist feature for all user line types.

```
<Sysname> system-view
```

```
[Sysname] password-control blacklist all-line
```

New feature: Configuring ND attack detection for a VSI

Configuring ND attack detection for a VSI

About this task

On a VXLAN network, you can configure user validity check for the associated VSI on the VTEP device. Different from the user validity check within a VLAN, which is performed on ND untrusted interfaces, the user validity check for a VSI is performed on ND untrusted ACs. The security entries used for user validity check and the check process are the same for a VLAN and a VSI.

Ethernet service instances that are associated with the VSI of a VXLAN are ACs. For more information, see *VXLAN Configuration Guide*.

Restrictions and guidelines

If both ND snooping and ND detection are enabled in a VSI, the device learns ND snooping entries only on the trusted ACs.

Procedure

1. Enter system view.
system-view
2. Enter VSI view.
vsi *vsi-name*
3. Enable ND attack detection.
ipv6 nd detection enable
By default, ND attack detection is disabled.
4. (Optional.) Configure the ND trusted AC.
 - a. Return to system view.
quit
 - b. Enter interface view.
interface *interface-type* *interface-number*
 - c. Enter Ethernet service instance view.
service-instance *instance-id*
 - d. Configure the AC as ND trusted AC.
ipv6 nd detection trust
By default, all ACs are ND untrusted ACs.

Command reference

Modified command: ipv6 nd detection enable

Syntax

```
ipv6 nd detection enable
undo ipv6 nd detection enable
```

Views

VLAN view
VSI view

Change description

Before modification: The command is not available in VSI view.

After modification: The command is available in VSI view.

Modified command: ipv6 nd detection trust

Syntax

```
ipv6 nd detection trust
undo ipv6 nd detection trust
```

Views

Layer 2 Ethernet interface view
Layer 2 aggregate interface view
Ethernet service instance view (AC)

Change description

Before modification: This command is not available in Ethernet service instance view (AC).

After modification: This command is available in Ethernet service instance view (AC).

Modified command: display ipv6 nd detection statistics

Syntax

```
display ipv6 nd detection statistics [ interface interface-type
interface-number [ service-instance instance-id ] ]
```

Views

Any view

Change description

Before modification: The command does not support the **service-instance** keyword. The command output does not support displaying AC information.

After modification: The **service-instance** keyword was added to the command. The command output supports displaying AC information.

Modified command: reset ipv6 nd detection statistics

Syntax

```
reset ipv6 nd detection statistics [ interface interface-type
interface-number [ service-instance instance-id ] ]
```

Views

User view

Change description

Before modification: The command does not support the **service-instance** keyword.

After modification: The **service-instance** keyword was added to the command.

New feature: MVXLAN

MVXLAN overview

Multicast VXLAN (MVXLAN) transmits multicast traffic from a multicast source to multicast receivers in an EVPN network that uses VXLAN in the data plane.

Restrictions: Underlay network restriction

MVXLAN supports only IPv4 underlay networks in the current software version.

MVXLAN benefits

MVXLAN provides the following benefits:

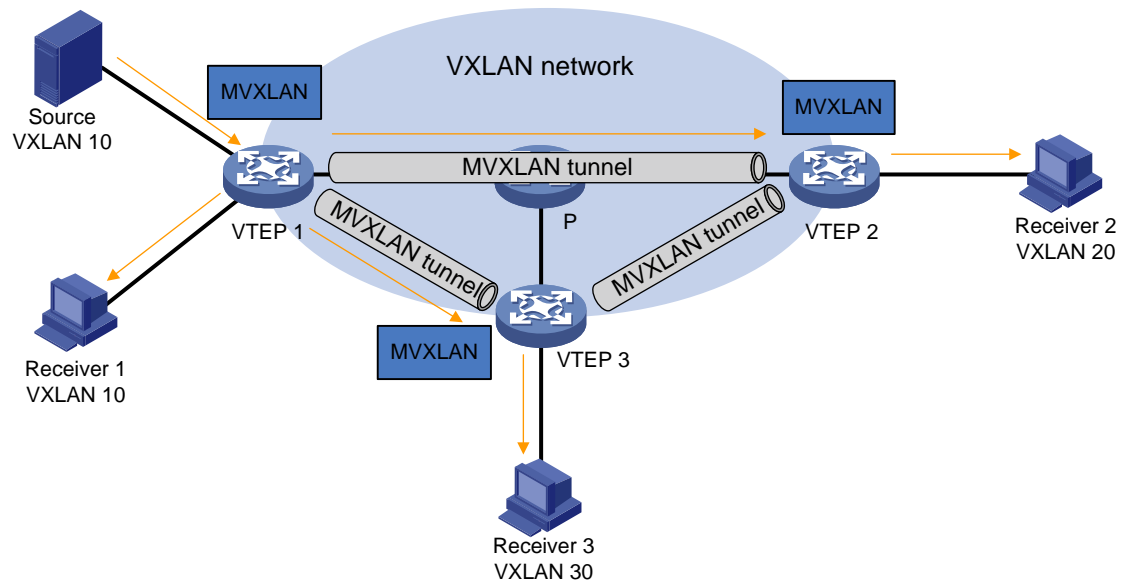
- **On-demand multicast forwarding**—Creates multicast distribution trees and manages multicast group members by using BGP EVPN routes and PIM.
- **Inter-VXLAN multicast forwarding**—Uses distributed EVPN gateways to forward Layer 3 multicast traffic between VXLANs.

MVXLAN network model

As shown in [Figure 2](#), distributed EVPN gateways are collocated with the VTEPs, and MVXLANs are created on the VTEPs to direct multicast traffic forwarding. When receiving multicast packets, a VTEP forwards them through ACs and MVXLAN tunnels to multicast receivers.

For more information about VTEPs, VSIs, and VXLANs, see *VXLAN Configuration Guide*. For more information about EVPN configuration, see "Configuring EVPN."

Figure 2 MVXLAN network model



Basic concepts in MVXLAN

The following are the basic concepts in MVXLAN:

- **Multicast distribution tree (MDT)**—An MDT is a multicast distribution tree constructed by all VTEPs in the same MVXLAN. MDTs include the default MDT and the data MDT.
- **Default group**—A default group is a unique multicast address assigned to each MVXLAN on the public network. It is the unique identifier of an MVXLAN on the public network and helps build the default MDT for an MVXLAN on the public network. Packets of the private multicast groups in an MVXLAN are encapsulated into packets of the default group before they are transmitted on the public network.
- **Default MDT**—A default MDT uses a default group address as its group address. The default MDT of an MVXLAN is uniquely identified by the default group and transmits all private multicast packets of the MVXLAN. A default MDT is automatically created after the default group is specified and will always exist on the public network, regardless of whether multicast services exist on the public network or MVXLAN.
- **Data group**—An MVXLAN is assigned a unique data group for MDT switchover. When the multicast traffic of an MVXLAN reaches or exceeds a threshold, the ingress VTEP selects a least used address from the data group range to encapsulate the multicast packets of the MVXLAN. Other VTEPs are notified to use the address to forward the traffic of the MVXLAN. This initiates the switchover to the data MDT.
- **Data MDT**—A data MDT is an MDT that uses a data group as its group address. At MDT switchover, VTEPs with downstream receivers join a data group to build a data MDT. The ingress VTEP forwards the encapsulated MVXLAN multicast traffic along the data MDT over the public network.

How MVXLAN works

MDT-based transmission

On the public network, MVXLAN multicast traffic is forwarded along an MDT rooted at the multicast source-side VTEP to leaf receiver-side VTEPs through unidirectional MVXLAN tunnels. MDT-based transmission ensures that multicast traffic is forwarded along optimal paths.

MP-BGP extension for MVXLAN

To support MVXLAN, MP-BGP introduces the following routes for creating MDTs under the EVPN address family:

- **Supplementary broadcast domain selective multicast Ethernet tag (SBD-SMET) route**—Contains private multicast source address and private multicast group address information. A receiver-side VTEP uses the SBD-SMET route to advertise its interest in a specific (*, G) or (S, G). An SBD-SMET route carries the RD configured in VPN instance view and export targets configured in VPN instance IPv4 address family view.
- **Selective provider multicast service interface route**—Also known as S-PMSI A-D route. An S-PMSI A-D route contains the private multicast source address, private multicast group address, default or data group address, and MVXLAN source interface address. S-PMSI A-D routes are used by the multicast source-side VTEP and its BGP peers to establish the default MDT and switch traffic from the default MDT to a data MDT. An S-PMSI A-D route carries the RD configured in VPN instance view and export targets configured in VPN instance IPv4 address family view.

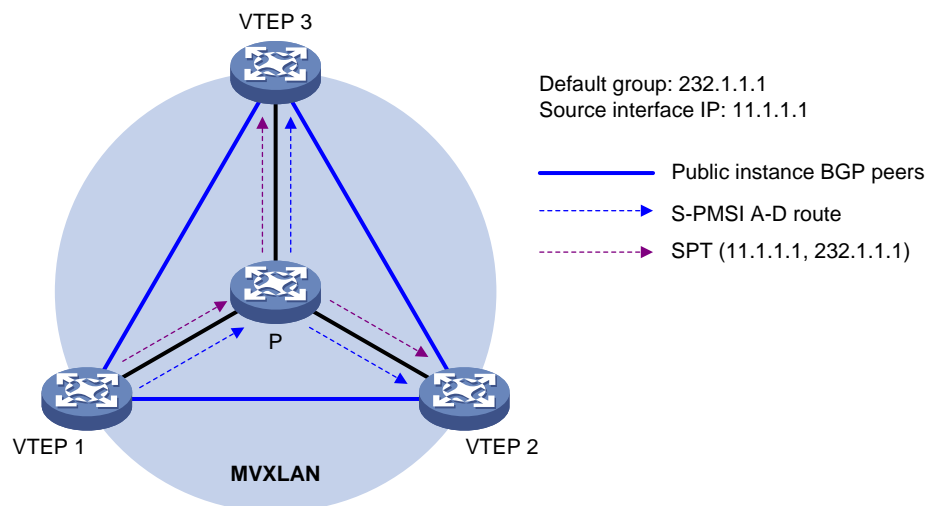
Automatic MVXLAN tunnel establishment and assignment

In an MVXLAN network, VTEPs automatically establish MVXLAN tunnels and assign them to MVXLANs to forward Layer 3 multicast traffic. The tunnel source is the MVXLAN source interface address, and the tunnel destination is the default or data group address. An MVXLAN tunnel is a unidirectional tunnel from the multicast source-side VTEP to a multicast receiver-side VTEP.

Default MDT establishment

The multicast routing protocol running on the public network can be PIM-SM or PIM-SSM. The process of creating a default MDT is the same in these PIM modes. All VTEPs in an MVXLAN join the default MDT of the MVXLAN. The private multicast packets of the MVXLAN are forwarded along the default MDT to the VTEPs, no matter whether the site attached to a VTEP contains receivers.

Figure 3 Default MDT establishment in a PIM-SM network



As shown in [Figure 3](#), PIM-SM runs on the public network, and MVXLAN is configured on all VTEPs. The process for establishing a default MDT is as follows:

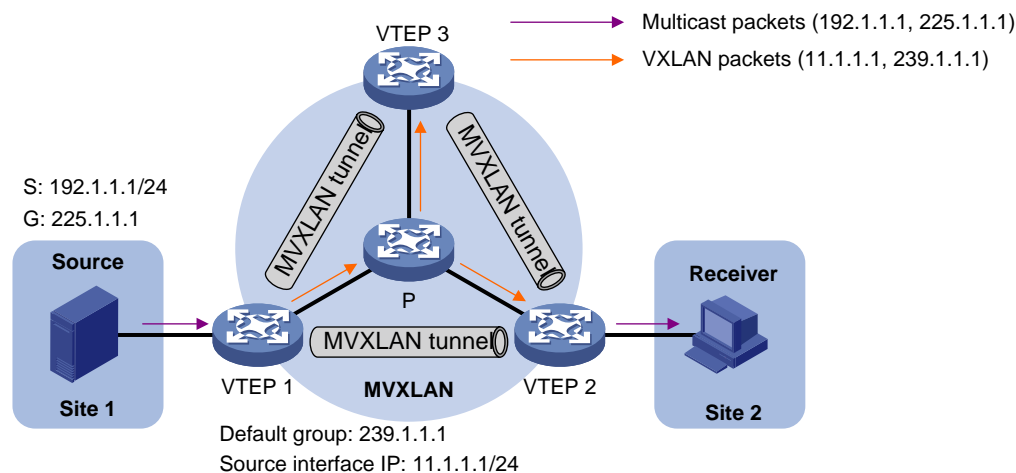
1. VTEP 1 sends an S-PMSI A-D route that contains (*, *) to VTEP 2 and VTEP 3.
2. VTEP 2 and VTEP 3 receive the route and join a multicast group according to the PMSI tunnel attribute of the route. The PMSI tunnel attribute contains the following information:

- The multicast source is the IP address of the MVXLAN source interface on VTEP 1.
 - The multicast group is the default group configured on VTEP 1.
3. Multicast forwarding entries are created on each device along the paths on the public network, and a shortest path tree (SPT) with VTEP 1 as the root and VTEP 2 and VTEP 3 as leaves is created. The SPT is the default MDT.

Default MDT-based transmission

After the default MDT is established, the multicast source sends the private multicast traffic to the receivers in each site along the default MDT. The private multicast packets are encapsulated into public multicast packets on the local VTEP and transmitted along the default MDT. Then, they are decapsulated on the remote VTEPs and transmitted in remote VXLAN sites.

Figure 2 Multicast data packet transmission



As shown in Figure 2, PIM-SM runs on the public network, the multicast source is attached to VTEP 1, and the multicast receiver is attached to VTEP 2. The multicast forwarding process is as follows:

1. The multicast source sends private multicast packets (192.1.1.1, 225.1.1.1) to VTEP 1.
2. VTEP 1 creates a multicast forwarding entry for (192.1.1.1, 225.1.1.1).
3. VTEP 1 processes the packets based on whether the receiver has joined the private multicast group:
 - If the receiver has sent an IGMP join message to VTEP 2, VTEP 1 has an SBD-SMET route sent by VTEP 1 that contains (*, G). VTEP 1 adds VXLAN encapsulation to the packets according to the route and forwards them to VTEP 2 and VTEP 3 along the default MDT. In the outer IP header of the VXLAN packets, the source IP address is the IP address of the MVXLAN source interface, and the destination IP address is the default group address.
 - If no receiver exists, VTEP 1 drops the packets.
4. VTEP 2 decapsulates the VXLAN packets and forwards the private multicast packets to the receiver.
5. VTEP 3 decapsulates the VXLAN packets and drops the private multicast packets because no local receiver exists.

MDT switchover

An MVXLAN can use the default MDT or a data MDT for multicast traffic forwarding. The default MDT is uniquely identified by the default group, and a data MDT is uniquely identified by a data group. Each default group is associated with a data group range.

Switching from the default MDT to a data MDT

When a multicast packet of an MVXLAN is transmitted through the default MDT on the public network, the packet is forwarded to all VTEPs configured with the VPN instance of the MVXLAN. This occurs whether or not any active receivers exist in the sites attached to the VTEPs. When the rate of the multicast traffic of that MVXLAN is high, multicast traffic might be flooded on the public network. This increases the bandwidth use and brings extra burden on the VTEPs.

To optimize multicast transmission, the MDT-based MVXLAN solution introduces a dedicated data MDT. The data MDT is built between the VTEPs that are attached to MVXLAN multicast receivers and multicast sources. When specific network criteria are met, multicast traffic is switched from the default MDT to the data MDT.

A switchover from the default MDT to the data MDT is initiated as follows:

1. The source-side VTEP periodically examines the forwarding rate of the MVXLAN multicast traffic. The default MDT switches to the data MDT only when both of the following criteria are met:
 - The private multicast traffic has passed the ACL rule filtering for default MDT to data MDT switchover.
 - The traffic rate of the private multicast stream has stayed above the switchover threshold for a certain period of time.
2. The source-side VTEP selects a least-used address from the data group range and sends an S-PMSI A-D route to all the other VTEPs down the default MDT. This route contains the private multicast source address, private multicast group address, IP address of the MVXLAN source interface, and data group address.
3. Each VTEP that receives the route examines whether it has receivers of that private multicast stream.
If so, it joins the data MDT rooted at the source-side VTEP. Otherwise, it caches the route and will join the data MDT when it has attached receivers.
4. After sending the S-PMSI A-D route, the source-side VTEP starts the data-delay timer. When the timer expires, the source-side VTEP uses the data group address to encapsulate the private multicast traffic. The multicast traffic is then forwarded down the data MDT.
5. After the multicast traffic is switched from the default MDT to the data MDT, a downstream VTEP can leave the data MDT by sending a PIM prune message if it no longer has active receivers attached to it.

Switching from the data MDT to the default MDT

After the MVXLAN multicast traffic is switched to the data MDT, the multicast traffic conditions might change and no longer meet the switchover criterion. In this case, the source-side VTEP initiates a backward MDT switchover process when any of the following criteria are met:

- The traffic rate of the MVXLAN multicast traffic has stayed below the switchover threshold for a certain period of time.
- The associated data group range is changed, and the data group address for encapsulating the MVXLAN multicast traffic is not in the new address range.
- The ACL rule for controlling the switchover from the default MDT to the data MDT has changed, and the MVXLAN multicast traffic fails to pass the new ACL rule.

MVXLAN tasks at a glance

To configure MVXLAN, perform the following tasks:

1. Configuring EVPN
 - a. Configuring a VXLAN on a VSI
 - b. Mapping ACs to a VSI
 - c. Configuring an EVPN instance

- d. Configuring BGP to advertise BGP EVPN routes
 - e. Configuring a distributed EVPN gateway

For more information about EVPN configuration, see "Configuring EVPN."
- 2. Configuring IGMP and IGMP snooping
 - a. Enabling IGMP on a VSI interface
 - b. Enabling IGMP snooping
 - c. Configuring IGMP snooping proxying

For more information about IGMP and IGMP snooping configuration, see *IP Multicast Configuration Guide*.
- 3. Configuring PIM on the transport-facing interfaces of VTEPs

Choose one of the following tasks:

 - o Configuring PIM-SM
 - o Configuring PIM-SSM

For more information about PIM configuration, see *IP Multicast Configuration Guide*.
- 4. Configuring MVXLAN
 - a. [Enabling IP multicast routing for a VPN instance](#)
 - b. [Creating an MVXLAN](#)
 - c. [Configuring a default group](#)
 - d. [Specifying the MVXLAN source interface](#)
 - e. [Configuring MDT switchover parameters](#)
 - f. [Configuring a VSI interface as a distributed DR interface](#)
 - g. [Configuring an MVXLAN extranet RPF selection policy](#)

Enabling IP multicast routing for a VPN instance

1. Enter system view.
system-view
2. Enable IP multicast routing for a VPN instance and enter MRIB view.
multicast routing vpn-instance *instance-name*
By default, IP multicast routing is disabled for VPN instances.
For more information about this command, see multicast routing and forwarding commands in *IP Multicast Command Reference*.

Creating an MVXLAN

About MVXLANS

You can create one or multiple MDT-based MVXLANS on a VTEP to provide services for different VPN instances and the public instance.

Creating an MVXLAN for a VPN instance

1. Enter system view.
system-view
2. Create an MDT-based MVXLAN and enter MVXLAN view.
multicast-vpn vxlan vpn-instance *instance-name* **mode mdt**

Creating an MVXLAN for the public instance

1. Enter system view.
system-view
2. Create an MDT-based MVXLAN and enter MVXLAN view.
multicast-vpn vxlan public-instance mode mdt

Configuring a default group

About a default group

When adding VXLAN encapsulation to private multicast packets, the VTEP uses the default group as the destination IP address in the outer IP header.

Restrictions and guidelines

The default group address of an MVXLAN must be unique among MVXLANs, and it cannot be the same as a data group address of any MVXLAN.

Procedure

1. Enter system view.
system-view
2. Enter MVXLAN view.
multicast-vpn vxlan vpn-instance *instance-name* mode mdt
3. Create the MVXLAN IPv4 address family and enter its view.
address-family ipv4
4. Configure the default group.
default-group *group-address*
By default, no default group exists.

Specifying the MVXLAN source interface

About the MVXLAN source interface

When adding VXLAN encapsulation to private multicast packets, the VTEP uses the IP address of the MVXLAN source interface as the source IP address in the outer IP header.

Restrictions and guidelines

For the VTEP to obtain correct routing information, you must specify the interface used for establishing BGP peer relationships as the MVXLAN source interface.

Procedure

1. Enter system view.
system-view
2. Enter MVXLAN view.
multicast-vpn vxlan vpn-instance *instance-name* mode mdt
3. Enter MVXLAN IPv4 address family view.
address-family ipv4
4. Specify the MVXLAN source interface.
source *interface-type interface-number*
By default, no MVXLAN source interface is specified.

Configuring MDT switchover parameters

About MDT switchover parameters

The traffic rate of the private multicast traffic might fluctuate around the MDT switchover threshold. To avoid frequent switching of multicast traffic between the default MDT and a data MDT, set the data-delay period. MDT switchover does not take place immediately after the multicast traffic rate exceeds the switchover threshold. It takes place after a data-delay period, during which the traffic rate must stay above the switchover threshold.

Restrictions and guidelines

On a VTEP, the data group range of an MVXLAN cannot include the default group or data groups of any other MVXLAN.

All VPN instances share the data group resources. As a best practice to avoid data group resource exhaustion, specify a reasonable data group range for a VPN instance.

Procedure

1. Enter system view.
system-view
2. Enter MVXLAN view.
multicast-vpn vxlan vpn-instance *instance-name* **mode mdt**
3. Enter MVXLAN IPv4 address family view.
address-family ipv4
4. Configure the data group range and the switchover criteria.
data-group *group-address* { *mask-length* | *mask* } [**acl** *acl-number* | **threshold** *threshold-value*] *

By default, no data group range exists, and the default MDT to data MDT switchover never occurs.
5. Set the data-delay period.
data-delay *delay*

By default, the data-delay period is 3 seconds.

Configuring a VSI interface as a distributed DR interface

About a distributed DR interface

On the VTEPs configured with MVXLAN, you must specify the VSI interfaces that act as distributed EVPN gateways as distributed designated router (DR) interfaces. This operation ensures that a distributed EVPN gateway can forward multicast traffic to the local site.

Procedure

1. Enter system view.
system-view
2. Enter VSI interface view.
interface vsi-interface *interface-number*
3. Configure the VSI interface as a distributed DR interface.
pim distributed-dr

A VSI instance is not a distributed DR interface.

Configuring an MVXLAN extranet RPF selection policy

About MVXLAN extranet RPF selection policies

MVXLAN extranet RPF routing policies are used for multicast transmission when multicast sources and receivers are located in different VPNs.

Restrictions and guidelines

The PIM mode in the source VPN instance and the receiver VPN instance must be the same. Only PIM-SM and PIM-SSM are supported.

In PIM-SM mode, if multiple multicast groups require inter-VPN transmission, configure a dedicated RP for these multicast groups and specify the RP as the multicast source in the RPF selection policy.

In PIM-SM mode, you can configure only one RPF selection policy for a multicast source and multicast group in a VPN instance.

To implement source-specific RPF selection in MVXLAN extranet, you must configure two MVXLAN extranet RPF routing policies as follows:

- In one policy, specify the address of the RP designated to the multicast group that requires inter-VPN multicast communication as the source address.
- In the other policy, specify the multicast source in the source VPN instance as the source address.

To implement source-and-group-specific RPF selection in MVXLAN extranet, you must configure two MVXLAN extranet RPF routing policies as follows:

- In one policy, specify the address of the RP designated to the multicast group as the source address, and specify the multicast group.
- In the other policy, specify the multicast source in the source VPN instance as the source address, and specify the multicast group.

If an MVXLAN extranet RPF selection policy with only the multicast group address specified is configured in the receiver VPN instance, the intra-VPN multicast traffic transmission will be interrupted.

Multicast source addresses in different MVXLAN extranet RPF routing policies cannot be the same, but they can overlap. The same restriction applies to the multicast group addresses in different MVXLAN extranet RPF routing policies. If multiple routing policies exist for an (S, G) entry, the device selects the policy in which the multicast group address has the longest mask. If multiple policies have the same mask length, the device selects the policy in which the multicast source address has the longest mask.

Procedure

1. Enter system view.
system-view
2. Enter MRIB view.
multicast routing [vpn-instance vpn-instance-name]
3. Configure an IPv4 MVXLAN extranet RPF selection policy.
multicast extranet select-rpf [vpn-instance vpn-instance-name]
{ source source-address { mask | mask-length } | group group-address
{ mask | mask-length } } *

Display and maintenance commands for MVXLAN

Execute **display** commands in any view.

| Task | Command |
|---|--|
| Display received data group information in an MVXLAN. | display multicast-vpn vxlan { vpn-instance instance-name public-instance } data-group receive [brief [active group group-address sender source-address vpn-source-address [mask { mask-length mask }] vpn-group-address [mask { mask-length mask }]] *] |
| Display sent data group information in an MVXLAN. | display multicast-vpn vxlan { vpn-instance instance-name public-instance } data-group send [group group-address vpn-source-address [mask { mask-length mask }] vpn-group-address [mask { mask-length mask }]] * |
| Display information about default groups. | display multicast-vpn vxlan [vpn-instance instance-name public-instance] default-group { local remote } |

MVXLAN configuration examples

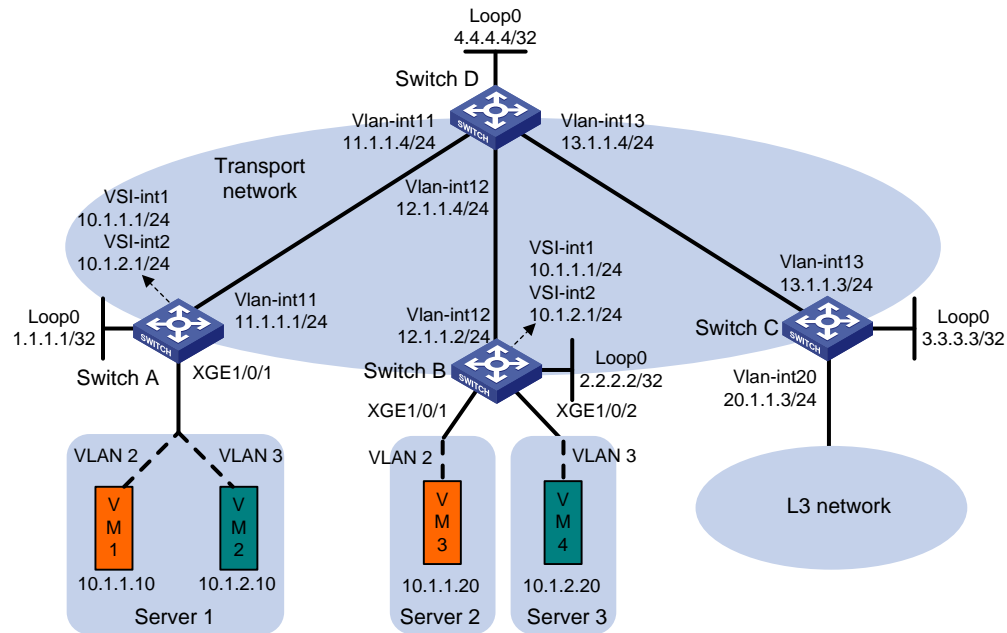
Example: Configuring inter-MVXLAN Layer 3 multicast forwarding

Network configuration

As shown in [Figure 2](#), VM 1 is the multicast source of multicast group 225.0.0.0, and the other VMs are multicast receivers. Configure MVXLAN to forward the multicast traffic from the source to the receivers.

- Configure VXLAN 10 and VXLAN 20 on Switch A and Switch B to extend VLAN 2 and VLAN 3 across the sites.
- Configure Switch A and Switch B as distributed EVPN gateways to provide gateway services. Configure Switch C as a border gateway to provide access to the connected Layer 3 network.
- Configure Switch D as an RR to reflect BGP EVPN routes between Switch A, Switch B, and Switch C.
- Configure PIM-SM on the transport-facing interfaces of Switches A through D. Configure IGMP snooping on Switches A through C for multicast forwarding entry creation.

Figure 2 Network diagram



Procedure

- On VM 1 and VM 3, specify 10.1.1.1 as the gateway address. On VM 2 and VM 4, specify 10.1.2.1 as the gateway address. (Details not shown.)
- Configure IP addresses and unicast routing settings:
 - # Assign IP addresses to interfaces, as shown in Figure 2. (Details not shown.)
 - # Configure OSPF on all transport network switches (Switches A through D) for them to reach one another. (Details not shown.)
- Configure Switch A:
 - # Enable L2VPN and IP multicast routing.


```
<SwitchA> system-view
[SwitchA] l2vpn enable
[SwitchA] multicast routing
[SwitchA-mrib] quit
```
 - # Enable the IGMP snooping feature.


```
[SwitchA] igmp-snooping
[SwitchA-igmp-snooping] quit
```
 - # Disable remote MAC address learning and remote ARP learning.


```
[SwitchA] vxlan tunnel mac-learning disable
[SwitchA] vxlan tunnel arp-learning disable
```
 - # Create VLAN-interface 11 and enter its view.


```
[SwitchA] vlan 11
[SwitchA-vlan11] quit
[SwitchA] interface vlan-interface 11
```
 - # Enable PIM-SM on VLAN-interface 11.


```
[SwitchA-Vlan-interface11] pim sm
[SwitchA-Vlan-interface11] quit
```
 - # Create an EVPN instance on VSI **vpna**.


```
[SwitchA] vsi vpna
```

```

[SwitchA-vsi-vpna] evpn encapsulation vxlan
[SwitchA-vsi-vpna-evpn-vxlan] quit
# Enable IGMP snooping and IGMP snooping proxying on VSI vpna.
[SwitchA-vsi-vpna] igmp-snooping enable
[SwitchA-vsi-vpna] igmp-snooping proxy enable
# Create VXLAN 10.
[SwitchA-vsi-vpna] vxlan 10
[SwitchA-vsi-vpna-vxlan-10] quit
[SwitchA-vsi-vpna] quit
# Create an EVPN instance on VSI vpb.
[SwitchA] vsi vpb
[SwitchA-vsi-vpb] evpn encapsulation vxlan
[SwitchA-vsi-vpb-evpn-vxlan] quit
# Enable IGMP snooping and IGMP snooping proxying on VSI vpb.
[SwitchA-vsi-vpb] igmp-snooping enable
[SwitchA-vsi-vpb] igmp-snooping proxy enable
# Create VXLAN 20.
[SwitchA-vsi-vpb] vxlan 20
[SwitchA-vsi-vpb-vxlan-20] quit
[SwitchA-vsi-vpb] quit
# Configure BGP to advertise BGP EVPN routes.
[SwitchA] bgp 200
[SwitchA-bgp-default] peer 4.4.4.4 as-number 200
[SwitchA-bgp-default] peer 4.4.4.4 connect-interface loopback 0
[SwitchA-bgp-default] address-family l2vpn evpn
[SwitchA-bgp-default-evpn] peer 4.4.4.4 enable
[SwitchA-bgp-default-evpn] quit
[SwitchA-bgp-default] quit
# On Ten-GigabitEthernet 1/0/1, create Ethernet service instance 1000 to match VLAN 2.
[SwitchA] interface ten-gigabitethernet 1/0/1
[SwitchA-Ten-GigabitEthernet1/0/1] service-instance 1000
[SwitchA-Ten-GigabitEthernet1/0/1-srv1000] encapsulation s-vid 2
# Map Ethernet service instance 1000 to VSI vpna.
[SwitchA-Ten-GigabitEthernet1/0/1-srv1000] xconnect vsi vpna
[SwitchA-Ten-GigabitEthernet1/0/1-srv1000] quit
# On Ten-GigabitEthernet 1/0/1, create Ethernet service instance 2000 to match VLAN 3.
[SwitchA-Ten-GigabitEthernet1/0/1] service-instance 2000
[SwitchA-Ten-GigabitEthernet1/0/1-srv2000] encapsulation s-vid 3
# Map Ethernet service instance 2000 to VSI vpb.
[SwitchA-Ten-GigabitEthernet1/0/1-srv2000] xconnect vsi vpb
[SwitchA-Ten-GigabitEthernet1/0/1-srv2000] quit
[SwitchA-Ten-GigabitEthernet1/0/1] quit
# Configure RD and route target settings for VPN instance vpna.
[SwitchA] ip vpn-instance vpna
[SwitchA-vpn-instance-vpna] route-distinguisher 1:1
[SwitchA-vpn-instance-vpna] address-family ipv4
[SwitchA-vpn-ipv4-vpna] vpn-target 2:2

```

```
[SwitchA-vpn-ipv4-vpna] quit
[SwitchA-vpn-instance-vpna] address-family evpn
[SwitchA-vpn-evpn-vpna] vpn-target 1:1
[SwitchA-vpn-evpn-vpna] quit
[SwitchA-vpn-instance-vpna] quit
```

Configure VSI-interface 1.

```
[SwitchA] interface vsi-interface 1
[SwitchA-Vsi-interface1] ip binding vpn-instance vpna
[SwitchA-Vsi-interface1] ip address 10.1.1.1 255.255.255.0
[SwitchA-Vsi-interface1] pim sm
[SwitchA-Vsi-interface1] pim distributed-dr
[SwitchA-Vsi-interface1] mac-address 1-1-1
[SwitchA-Vsi-interface1] distributed-gateway local
[SwitchA-Vsi-interface1] local-proxy-arp enable
[SwitchA-Vsi-interface1] quit
```

Configure VSI-interface 2.

```
[SwitchA] interface vsi-interface 2
[SwitchA-Vsi-interface2] ip binding vpn-instance vpna
[SwitchA-Vsi-interface2] ip address 10.1.2.1 255.255.255.0
[SwitchA-Vsi-interface2] pim sm
[SwitchA-Vsi-interface2] pim distributed-dr
[SwitchA-Vsi-interface2] mac-address 2-2-2
[SwitchA-Vsi-interface2] distributed-gateway local
[SwitchA-Vsi-interface2] local-proxy-arp enable
[SwitchA-Vsi-interface2] quit
```

Associate VSI-interface 3 with VPN instance **vpna**, and configure the L3 VXLAN ID as 1000 for the VPN instance.

```
[SwitchA] interface vsi-interface 3
[SwitchA-Vsi-interface3] ip binding vpn-instance vpna
[SwitchA-Vsi-interface3] l3-vni 1000
[SwitchA-Vsi-interface3] pim sm
[SwitchA-Vsi-interface3] quit
```

Enable IP multicast routing for VPN instance **vpna**.

```
[SwitchA] multicast routing vpn-instance vpna
[SwitchA-mrib-vpna] quit
```

Create an MDT-based MVXLAN for VPN instance **vpna** and enter MVXLAN IPv4 address family view. Configure the default group, MVXLAN source interface, and data group range settings.

```
[SwitchA] multicast-vpn vxlan vpn-instance vpna mode mdt
[SwitchA-mvxlan-vpna] address-family ipv4
[SwitchA-mvxlan-vpna-ipv4] default-group 236.0.0.1
[SwitchA-mvxlan-vpna-ipv4] source loopback 0
[SwitchA-mvxlan-vpna-ipv4] data-group 239.0.1.0 24
[SwitchA-mvxlan-vpna-ipv4] quit
[SwitchA-mvxlan-vpna] quit
```

Configure Loopback 1.

```
[SwitchA] interface loopback 1
[SwitchA-LoopBack1] ip binding vpn-instance vpna
```

```
[SwitchA-LoopBack1] ip address 12.12.12.12 32
[SwitchA-LoopBack1] pim sm
[SwitchA-LoopBack1] quit
```

Enter VPN instance PIM view, and configure Loopback 1 as a candidate-BSR and candidate-RP in VPN instance **vpna.**

```
[SwitchA] pim vpn-instance vpna
[SwitchA-pim-vpna] c-bsr 12.12.12.12
[SwitchA-pim-vpna] c-rp 12.12.12.12
[SwitchA-pim-vpna] quit
```

Specify VSI-interface 1 as the gateway interface for VSI **vpna.**

```
[SwitchA] vsi vpna
[SwitchA-vsi-vpna] gateway vsi-interface 1
[SwitchA-vsi-vpna] quit
```

Specify VSI-interface 2 as the gateway interface for VSI **vpnb.**

```
[SwitchA] vsi vpb
[SwitchA-vsi-vpb] gateway vsi-interface 2
[SwitchA-vsi-vpb] quit
```

4. Configure Switch B:

Enable L2VPN and IP multicast routing.

```
<SwitchB> system-view
[SwitchB] l2vpn enable
[SwitchB] multicast routing
[SwitchB-mrib] quit
```

Enable the IGMP snooping feature.

```
[SwitchB] igmp-snooping
[SwitchB-igmp-snooping] quit
```

Disable remote MAC address learning and remote ARP learning.

```
[SwitchB] vxlan tunnel mac-learning disable
[SwitchB] vxlan tunnel arp-learning disable
```

Create VLAN-interface 12 and enter its view.

```
[SwitchB] vlan 12
[SwitchB-vlan12] quit
[SwitchB] interface vlan-interface 12
```

Enable PIM-SM on VLAN-interface 12.

```
[SwitchB-Vlan-interface12] pim sm
[SwitchB-Vlan-interface12] quit
```

Create an EVPN instance on VSI **vpna.**

```
[SwitchB] vsi vpna
[SwitchB-vsi-vpna] evpn encapsulation vxlan
[SwitchB-vsi-vpna-evpn-vxlan] quit
```

Enable IGMP snooping and IGMP snooping proxying on VSI **vpna.**

```
[SwitchB-vsi-vpna] igmp-snooping enable
[SwitchB-vsi-vpna] igmp-snooping proxy enable
```

Create VXLAN 10.

```
[SwitchB-vsi-vpna] vxlan 10
[SwitchB-vsi-vpna-vxlan-10] quit
[SwitchB-vsi-vpna] quit
```

Create an EVPN instance on VSI `vpnb`.

```
[SwitchB] vsi vpb  
[SwitchB-vsi-vpb] evpn encapsulation vxlan  
[SwitchB-vsi-vpb-evpn-vxlan] quit
```

Enable IGMP snooping and IGMP snooping proxying on VSI `vpnb`.

```
[SwitchB-vsi-vpb] igmp-snooping enable  
[SwitchB-vsi-vpb] igmp-snooping proxy enable
```

Create VXLAN 20.

```
[SwitchB-vsi-vpb] vxlan 20  
[SwitchB-vsi-vpb-vxlan-20] quit  
[SwitchB-vsi-vpb] quit
```

Configure BGP to advertise BGP EVPN routes.

```
[SwitchB] bgp 200  
[SwitchB-bgp-default] peer 4.4.4.4 as-number 200  
[SwitchB-bgp-default] peer 4.4.4.4 connect-interface loopback 0  
[SwitchB-bgp-default] address-family l2vpn evpn  
[SwitchB-bgp-default-evpn] peer 4.4.4.4 enable  
[SwitchB-bgp-default-evpn] quit  
[SwitchB-bgp-default] quit
```

On Ten-GigabitEthernet 1/0/1, create Ethernet service instance 1000 to match VLAN 2.

```
[SwitchB] interface ten-gigabitethernet 1/0/1  
[SwitchB-Ten-GigabitEthernet1/0/1] service-instance 1000  
[SwitchB-Ten-GigabitEthernet1/0/1-srv1000] encapsulation s-vid 2
```

Map Ethernet service instance 1000 to VSI `vpna`.

```
[SwitchB-Ten-GigabitEthernet1/0/1-srv1000] xconnect vsi vpna  
[SwitchB-Ten-GigabitEthernet1/0/1-srv1000] quit
```

On Ten-GigabitEthernet 1/0/1, create Ethernet service instance 2000 to match VLAN 3.

```
[SwitchB-Ten-GigabitEthernet1/0/1] service-instance 2000  
[SwitchB-Ten-GigabitEthernet1/0/1-srv2000] encapsulation s-vid 3
```

Map Ethernet service instance 2000 to VSI `vpnb`.

```
[SwitchB-Ten-GigabitEthernet1/0/1-srv2000] xconnect vsi vpb  
[SwitchB-Ten-GigabitEthernet1/0/1-srv2000] quit  
[SwitchB-Ten-GigabitEthernet1/0/1] quit
```

Configure RD and route target settings for VPN instance `vpna`.

```
[SwitchB] ip vpn-instance vpna  
[SwitchB-vpn-instance-vpna] route-distinguisher 1:1  
[SwitchB-vpn-instance-vpna] address-family ipv4  
[SwitchB-vpn-ipv4-vpna] vpn-target 2:2  
[SwitchB-vpn-ipv4-vpna] quit  
[SwitchB-vpn-instance-vpna] address-family evpn  
[SwitchB-vpn-evpn-vpna] vpn-target 1:1  
[SwitchB-vpn-evpn-vpna] quit  
[SwitchB-vpn-instance-vpna] quit
```

Configure VSI-interface 1.

```
[SwitchB] interface vsi-interface 1  
[SwitchB-Vsi-interfacel] ip binding vpn-instance vpna  
[SwitchB-Vsi-interfacel] ip address 10.1.1.1 255.255.255.0  
[SwitchB-Vsi-interfacel] igmp enable
```

```
[SwitchB-Vsi-interface1] mac-address 1-1-1
[SwitchB-Vsi-interface1] distributed-gateway local
[SwitchB-Vsi-interface1] local-proxy-arp enable
[SwitchB-Vsi-interface1] quit
```

Configure VSI-interface 2.

```
[SwitchB] interface vsi-interface 2
[SwitchB-Vsi-interface2] ip binding vpn-instance vpna
[SwitchB-Vsi-interface2] ip address 10.1.2.1 255.255.255.0
[SwitchB-Vsi-interface2] igmp enable
[SwitchB-Vsi-interface2] mac-address 2-2-2
[SwitchB-Vsi-interface2] distributed-gateway local
[SwitchB-Vsi-interface2] local-proxy-arp enable
[SwitchB-Vsi-interface2] quit
```

Associate VSI-interface 3 with VPN instance **vpna**, and configure the L3 VXLAN ID as 1000 for the VPN instance.

```
[SwitchB] interface vsi-interface 3
[SwitchB-Vsi-interface3] ip binding vpn-instance vpna
[SwitchB-Vsi-interface3] l3-vni 1000
[SwitchB-Vsi-interface3] pim sm
[SwitchB-Vsi-interface3] quit
```

Enable IP multicast routing for VPN instance **vpna**.

```
[SwitchB] multicast routing vpn-instance vpna
[SwitchB-mrib-vpna] quit
```

Create an MDT-based MVXLAN for VPN instance **vpna** and enter MVXLAN IPv4 address family view. Configure the MVXLAN source interface.

```
[SwitchB] multicast-vpn vxlan vpn-instance vpna mode mdt
[SwitchB-mvxlan-vpna] address-family ipv4
[SwitchB-mvxlan-vpna-ipv4] source loopback 0
[SwitchB-mvxlan-vpna-ipv4] quit
[SwitchB-mvxlan-vpna] quit
```

Configure Loopback 1.

```
[SwitchB] interface loopback 1
[SwitchB-LoopBack1] ip binding vpn-instance vpna
[SwitchB-LoopBack1] ip address 12.12.12.12 32
[SwitchB-LoopBack1] pim sm
[SwitchB-LoopBack1] quit
```

Enter VPN instance PIM view, and configure Loopback 1 as a candidate-BSR and candidate-RP in VPN instance **vpna**.

```
[SwitchB] pim vpn-instance vpna
[SwitchB-pim-vpna] c-bsr 12.12.12.12
[SwitchB-pim-vpna] c-rp 12.12.12.12
[SwitchB-pim-vpna] quit
```

Specify VSI-interface 1 as the gateway interface for VSI **vpna**.

```
[SwitchB] vsi vpna
[SwitchB-vsi-vpna] gateway vsi-interface 1
[SwitchB-vsi-vpna] quit
```

Specify VSI-interface 2 as the gateway interface for VSI **vpnb**.

```
[SwitchB] vsi vpb
```

```
[SwitchB-vsi-vpnb] gateway vsi-interface 2
[SwitchB-vsi-vpnb] quit
```

5. Configure Switch C:

Enable L2VPN and IP multicast routing.

```
<SwitchC> system-view
[SwitchC] l2vpn enable
[SwitchC] multicast routing
[SwitchC-mrib] quit
```

Disable remote MAC address learning and remote ARP learning.

```
[SwitchC] vxlan tunnel mac-learning disable
[SwitchC] vxlan tunnel arp-learning disable
```

Create VLAN-interface 13 and enter its view.

```
[SwitchC] vlan 13
[SwitchC-vlan13] quit
[SwitchC] interface vlan-interface 13
```

Enable PIM-SM on VLAN-interface 13.

```
[SwitchC-Vlan-interface13] pim sm
[SwitchC-Vlan-interface13] quit
```

Configure BGP to advertise BGP EVPN routes.

```
[SwitchC] bgp 200
[SwitchC-bgp-default] peer 4.4.4.4 as-number 200
[SwitchC-bgp-default] peer 4.4.4.4 connect-interface loopback 0
[SwitchC-bgp-default] address-family l2vpn evpn
[SwitchC-bgp-default-evpn] peer 4.4.4.4 enable
[SwitchC-bgp-default-evpn] quit
[SwitchC-bgp-default] quit
```

Configure RD and route target settings for VPN instance **vpna.**

```
[SwitchC] ip vpn-instance vpna
[SwitchC-vpn-instance-vpna] route-distinguisher 1:1
[SwitchC-vpn-instance-vpna] address-family ipv4
[SwitchC-vpn-ipv4-vpna] vpn-target 2:2
[SwitchC-vpn-ipv4-vpna] quit
[SwitchC-vpn-instance-vpna] address-family evpn
[SwitchC-vpn-evpn-vpna] vpn-target 1:1
[SwitchC-vpn-evpn-vpna] quit
[SwitchC-vpn-instance-vpna] quit
```

Associate VSI-interface 3 with VPN instance **vpna, and configure the L3 VXLAN ID as 1000 for the VPN instance.**

```
[SwitchC] interface vsi-interface 3
[SwitchC-Vsi-interface3] ip binding vpn-instance vpna
[SwitchC-Vsi-interface3] l3-vni 1000
[SwitchC-Vsi-interface3] pim sm
[SwitchC-Vsi-interface3] quit
```

Enable IP multicast routing on VPN instance **vpna.**

```
[SwitchC] multicast routing vpn-instance vpna
[SwitchC-mrib-vpna] quit
```

Create an MDT-based MVXLAN for VPN instance **vpna and enter MVXLAN IPv4 address family view. Configure the MVXLAN source interface.**

```
[SwitchC] multicast-vpn vxlan vpn-instance vpna mode mdt
[SwitchC-mvxlan-vpna] address-family ipv4
[SwitchC-mvxlan-vpna-ipv4] source loopback 0
[SwitchC-mvxlan-vpna-ipv4] quit
[SwitchC-mvxlan-vpna] quit
```

Configure Loopback 1.

```
[SwitchC] interface loopback 1
[SwitchC-LoopBack1] ip binding vpn-instance vpna
[SwitchC-LoopBack1] ip address 12.12.12.12 32
[SwitchC-LoopBack1] pim sm
[SwitchC-LoopBack1] quit
```

Enter VPN instance PIM view, and configure Loopback 1 as a candidate-BSR and candidate-RP in VPN instance **vpna**.

```
[SwitchC] pim vpn-instance vpna
[SwitchC-pim-vpna] c-bsr 12.12.12.12
[SwitchC-pim-vpna] c-rp 12.12.12.12
[SwitchC-pim-vpna] quit
```

Configure a default route. Specify the next hop as 20.1.1.100, the IP address of a device in the Layer 3 network.

```
[SwitchC] ip route-static vpn-instance vpna 0.0.0.0 0 20.1.1.100
```

Import the default route to the BGP IPv4 unicast routing table of VPN instance **vpna**.

```
[SwitchC] bgp 200
[SwitchC-bgp-default] ip vpn-instance vpna
[SwitchC-bgp-default-vpna] address-family ipv4 unicast
[SwitchC-bgp-default-ipv4-vpna] default-route imported
[SwitchC-bgp-default-ipv4-vpna] import-route static
[SwitchC-bgp-default-ipv4-vpna] quit
[SwitchC-bgp-default-vpna] quit
[SwitchC-bgp-default] quit
```

Associate Ten-GigabitEthernet 1/0/2 with VPN instance **vpna**.

```
[SwitchC] interface ten-gigabitethernet 1/0/2
[SwitchC-Ten-GigabitEthernet1/0/2] ip binding vpn-instance vpna
[SwitchC-Ten-GigabitEthernet1/0/2] ip address 20.1.1.3 24
[SwitchC-Ten-GigabitEthernet1/0/2] pim sm
[SwitchC-Ten-GigabitEthernet1/0/2] quit
```

6. Configure Switch D:

Enable IP multicast routing.

```
<SwitchD> system-view
[SwitchD] multicast routing
[SwitchD-mrib] quit
```

Enter PIM view, and configure Loopback 0 as a candidate-BSR and candidate-RP in the public network.

```
[SwitchD] pim
[SwitchD-pim] c-bsr 4.4.4.4
[SwitchD] c-rp 4.4.4.4
[SwitchD] quit
```

Enable PIM-SM on VLAN-interface 11.

```
[SwitchD] interface vlan-interface 11
```

```

[SwitchD-Vlan-interface11] pim sm
[SwitchD-Vlan-interface11] quit
# Enable PIM-SM on VLAN-interface 12.
[SwitchD] interface vlan-interface12
[SwitchD-Vlan-interface12] pim sm
[SwitchD-Vlan-interface12] quit
# Enable PIM-SM on VLAN-interface 13.
[SwitchD] interface vlan-interface13
[SwitchD-Vlan-interface13] pim sm
[SwitchD-Vlan-interface13] quit
# Establish BGP connections with other transport network switches.
[SwitchD] bgp 200
[SwitchD-bgp-default] group evpn
[SwitchD-bgp-default] peer 1.1.1.1 group evpn
[SwitchD-bgp-default] peer 2.2.2.2 group evpn
[SwitchD-bgp-default] peer 3.3.3.3 group evpn
[SwitchD-bgp-default] peer evpn as-number 200
[SwitchD-bgp-default] peer evpn connect-interface loopback 0
# Configure BGP to advertise BGP EVPN routes, and disable route target filtering of received BGP EVPN routes.
[SwitchD-bgp-default] address-family l2vpn evpn
[SwitchD-bgp-default-evpn] peer evpn enable
[SwitchD-bgp-default-evpn] undo policy vpn-target
# Configure Switch D as an RR.
[SwitchD-bgp-default-evpn] peer evpn reflect-client
[SwitchD-bgp-default-evpn] quit
[SwitchD-bgp-default] quit

```

Verifying the configuration

1. Verify the multicast routing information on Switch A:

Verify that Switch A has multicast routing entries for VPN instance **vpna**.

```

<SwitchA> display pim vpn-instance vpna routing-table
Total 1 (*, G) entries; 1 (S, G) entries
(*, 225.0.0.0)
  RP: 12.12.12.12 (local)
  Protocol: pim-sm, Flag: WC RC
  UpTime: 02:57:31
  Upstream interface: Register-Tunnel0
    Upstream neighbor: NULL
    RPF prime neighbor: NULL
  Downstream interface information:
  Total number of downstream interfaces: 1
    1: MTunnel0
      Protocol: MD, UpTime: 02:57:31, Expires: -

(10.1.1.10, 225.0.0.0)
  RP: 12.12.12.12 (local)
  Protocol: pim-sm, Flag: SPT 2MSDP LOC ACT SQ RC 2MVPN

```

```

UpTime: 04:44:08
Upstream interface: Vsi-interfacel
  Upstream neighbor: NULL
  RPF prime neighbor: NULL
Downstream interface information:
Total number of downstream interfaces: 1
  1: MTunnell
    Protocol: MD, UpTime: 02:00:27, Expires: -

```

Verify that Switch A has multicast routing entries for the public network.

```

<SwitchA> display pim routing-table
Total 0 (*, G) entries; 2 (S, G) entries

```

```

(1.1.1.1, 236.0.0.1)
  RP: 4.4.4.4
  Protocol: pim-sm, Flag: SPT LOC VXLAN_L3
  UpTime: 02:09:52
  Upstream interface: MTunnel0 (VPN: vpna)
    Upstream neighbor: NULL
    RPF prime neighbor: NULL
  Downstream interface information:
Total number of downstream interfaces: 1
  1: Vlan-interfacell
    Protocol: pim-sm, UpTime: 01:16:34, Expires: 00:03:10

```

```

(1.1.1.1, 239.0.1.0)
  RP: 4.4.4.4
  Protocol: pim-sm, Flag: SPT LOC VXLAN_L3
  UpTime: 02:08:52
  Upstream interface: MTunnell1 (VPN: vpna)
    Upstream neighbor: NULL
    RPF prime neighbor: NULL
  Downstream interface information:
Total number of downstream interfaces: 1
  1: Vlan-interfacell
    Protocol: pim-sm, UpTime: 01:15:34, Expires: 00:03:11

```

2. Verify the multicast routing information on Switch B:

Verify that Switch B has multicast routing entries for VPN instance vpna.

```

<SwitchB> display pim vpn-instance vpna routing-table
Total 1 (*, G) entries; 1 (S, G) entries
(*, 225.0.0.0)
  RP: 12.12.12.12 (local)
  Protocol: pim-sm, Flag: WC
  UpTime: 05:04:06
  Upstream interface: Register-Tunnel0
    Upstream neighbor: NULL
    RPF prime neighbor: NULL
  Downstream interface information:
Total number of downstream interfaces: 1

```

```

1: Vsi-interface1
    Protocol: igmp, UpTime: 05:04:06, Expires: -

(10.1.1.10, 225.0.0.0)
    RP: 12.12.12.12 (local)
    Protocol: pim-sm, Flag: SPT ACT RQ FRTNL
    UpTime: 01:57:12
    Upstream interface: Multicast-UPE0 (0.0.0.0)
        Upstream neighbor: NULL
        RPF prime neighbor: NULL
    Downstream interface information:
    Total number of downstream interfaces: 1
        1: Vsi-interface1
            Protocol: pim-sm, UpTime: 01:57:12, Expires: -

# Verify that Switch B has multicast routing entries for the public network.
<SwitchB> display pim routing-table
Total 0 (*, G) entries; 2 (S, G) entries
(1.1.1.1, 236.0.0.1)
    RP: 4.4.4.4
    Protocol: pim-sm, Flag: SPT
    UpTime: 01:59:46
    Upstream interface: Vlan-interface12
        Upstream neighbor: 12.1.1.4
        RPF prime neighbor: 12.1.1.4
    Downstream interface information:
    Total number of downstream interfaces: 1
        1: Multicast-UPE0
            Protocol: MD, UpTime: 01:59:46, Expires: -

(1.1.1.1, 239.0.1.0)
    RP: 4.4.4.4
    Protocol: pim-sm, Flag: SPT ACT
    UpTime: 01:58:46
    Upstream interface: Vlan-interface12
        Upstream neighbor: 12.1.1.4
        RPF prime neighbor: 12.1.1.4
    Downstream interface information:
    Total number of downstream interfaces: 1
        1: Multicast-UPE0
            Protocol: MD, UpTime: 01:58:46, Expires: -

```

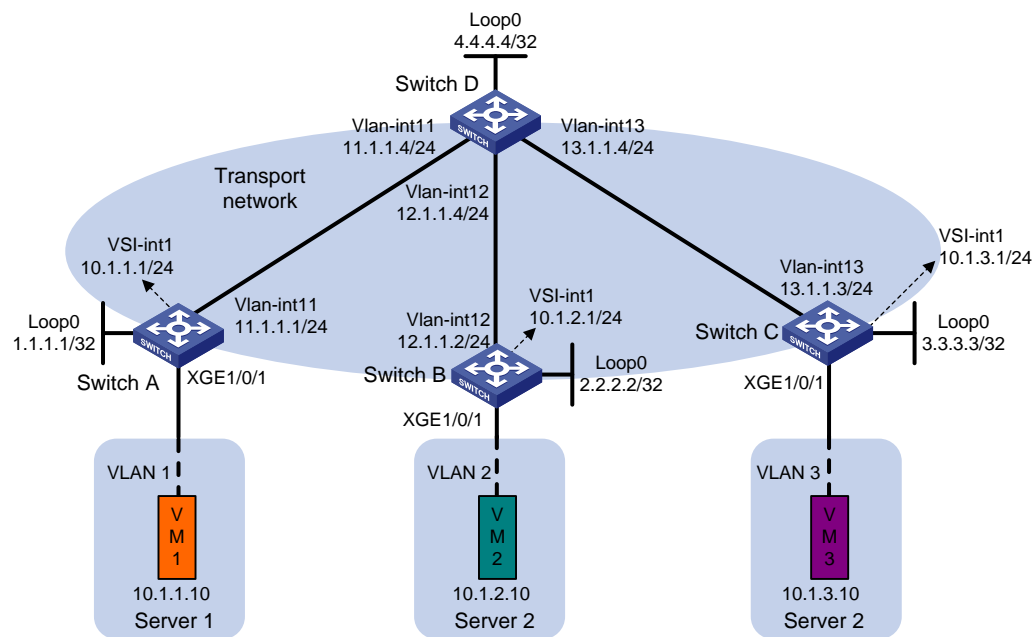
Example: Configuring MVXLAN Layer 3 multicast forwarding between private and public networks

Network configuration

As shown in [Figure 2](#), VM 1 is the multicast source, and the other VMs are multicast receivers. Configure MVXLAN to forward the multicast traffic from the source to the receivers.

- Configure VXLAN 10, VXLAN 20, and VXLAN 30 on Switch A, Switch B, and Switch C to meet the following requirements:
 - VXLAN 10 and VXLAN 20 are on the private network, and VXLAN 30 is on the public network.
 - VXLAN 10 can communicate with VXLAN 20 and VXLAN 30, and VXLAN 20 is isolated from VXLAN 30.
- Configure Switch A, Switch B, and Switch C as distributed EVPN gateways to provide gateway services for the VXLANs.
- Configure Switch D as an RR to reflect BGP EVPN routes between Switch A, Switch B, and Switch C.
- Configure PIM-SM on the transport-facing interfaces of Switches A through D. Configure IGMP snooping on Switches A through C for multicast forwarding entry creation.

Figure 2 Network diagram



Procedure

1. On VM 1, VM 2, and VM 3, specify 10.1.1.1, 10.1.2.1, and 10.1.3.1 as the gateway address, respectively. (Details not shown.)
2. Configure IP addresses and unicast routing settings:
 - # Assign IP addresses to interfaces, as shown in [Figure 2](#). (Details not shown.)
 - # Configure OSPF on all transport network switches (Switches A through D) for them to reach one another. (Details not shown.)
3. Configure Switch A:
 - # Enable L2VPN and IP multicast routing.

```
<SwitchA> system-view
[SwitchA] l2vpn enable
[SwitchA] multicast routing
[SwitchA-mrib] quit
```

 - # Enable the IGMP snooping feature.

```
[SwitchA] igmp-snooping
[SwitchA-igmp-snooping] quit
```

```

# Disable remote MAC address learning and remote ARP learning.
[SwitchA] vxlan tunnel mac-learning disable
[SwitchA] vxlan tunnel arp-learning disable
# Create VLAN-interface 11 and enter its view.
[SwitchA] vlan 11
[SwitchA-vlan11] quit
[SwitchA] interface vlan-interface 11
# Enable PIM-SM on VLAN-interface 11.
[SwitchA-Vlan-interface11] pim sm
[SwitchA-Vlan-interface11] quit
# Create an EVPN instance on VSI vpna.
[SwitchA] vsi vpna
[SwitchA-vsi-vpna] evpn encapsulation vxlan
[SwitchA-vsi-vpna-evpn-vxlan] quit
# Enable IGMP snooping and IGMP snooping proxying on VSI vpna.
[SwitchA-vsi-vpna] igmp-snooping enable
[SwitchA-vsi-vpna] igmp-snooping proxy enable
# Create VXLAN 10.
[SwitchA-vsi-vpna] vxlan 10
[SwitchA-vsi-vpna-vxlan-10] quit
[SwitchA-vsi-vpna] quit
# Configure BGP to advertise BGP EVPN routes.
[SwitchA] bgp 200
[SwitchA-bgp-default] peer 4.4.4.4 as-number 200
[SwitchA-bgp-default] peer 4.4.4.4 connect-interface loopback 0
[SwitchA-bgp-default] address-family l2vpn evpn
[SwitchA-bgp-default-evpn] peer 4.4.4.4 enable
[SwitchA-bgp-default-evpn] quit
[SwitchA-bgp-default] quit
# On Ten-GigabitEthernet 1/0/1, create Ethernet service instance 1000 to match VLAN 1.
[SwitchA] interface ten-gigabitethernet 1/0/1
[SwitchA-Ten-GigabitEthernet1/0/1] service-instance 1000
[SwitchA-Ten-GigabitEthernet1/0/1-srv1000] encapsulation s-vid 1
# Map Ethernet service instance 1000 to VSI vpna.
[SwitchA-Ten-GigabitEthernet1/0/1-srv1000] xconnect vsi vpna
[SwitchA-Ten-GigabitEthernet1/0/1-srv1000] quit
[SwitchA-Ten-GigabitEthernet1/0/1] quit
# Configure RD and route target settings for VPN instance vpna.
[SwitchA] ip vpn-instance vpna
[SwitchA-vpn-instance-vpna] route-distinguisher 1:1
[SwitchA-vpn-instance-vpna] address-family ipv4
[SwitchA-vpn-ipv4-vpna] vpn-target 2:2
[SwitchA-vpn-ipv4-vpna] quit
[SwitchA-vpn-instance-vpna] address-family evpn
[SwitchA-vpn-evpn-vpna] vpn-target 1:1
[SwitchA-vpn-evpn-vpna] quit
[SwitchA-vpn-instance-vpna] quit

```

Configure VSI-interface 1.

```
[SwitchA] interface vsi-interface 1
[SwitchA-Vsi-interface1] ip binding vpn-instance vpna
[SwitchA-Vsi-interface1] ip address 10.1.1.1 255.255.255.0
[SwitchA-Vsi-interface1] pim sm
[SwitchA-Vsi-interface1] pim distributed-dr
[SwitchA-Vsi-interface1] mac-address 1-1-1
[SwitchA-Vsi-interface1] distributed-gateway local
[SwitchA-Vsi-interface1] local-proxy-arp enable
[SwitchA-Vsi-interface1] quit
```

Associate VSI-interface 2 with VPN instance **vpna**, and configure the L3 VXLAN ID as 1000 for the VPN instance.

```
[SwitchA] interface vsi-interface 2
[SwitchA-Vsi-interface2] ip binding vpn-instance vpna
[SwitchA-Vsi-interface2] l3-vni 1000
[SwitchA-Vsi-interface2] pim sm
[SwitchA-Vsi-interface2] quit
```

Enable IP multicast routing on VPN instance **vpna**.

```
[SwitchA] multicast routing vpn-instance vpna
[SwitchA-mrib-vpna] quit
```

Create an MDT-based MVXLAN for VPN instance **vpna** and enter MVXLAN IPv4 address family view. Configure the default group, MVXLAN source interface, and data group range settings.

```
[SwitchA] multicast-vpn vxlan vpn-instance vpna mode mdt
[SwitchA-mvxlan-vpna] address-family ipv4
[SwitchA-mvxlan-vpna-ipv4] default-group 236.2.0.0
[SwitchA-mvxlan-vpna-ipv4] source loopback 0
[SwitchA-mvxlan-vpna-ipv4] data-group 239.2.0.0 24
[SwitchA-mvxlan-vpna-ipv4] quit
[SwitchA-mvxlan-vpna] quit
```

Configure Loopback 1.

```
[SwitchA] interface loopback 1
[SwitchA-LoopBack1] ip binding vpn-instance vpna
[SwitchA-LoopBack1] ip address 11.11.11.11 32
[SwitchA-LoopBack1] pim sm
[SwitchA-LoopBack1] quit
```

Enter VPN instance PIM view, and configure Loopback 1 as a candidate-BSR and candidate-RP in VPN instance **vpna**.

```
[SwitchA] pim vpn-instance vpna
[SwitchA-pim-vpna] c-bsr 11.11.11.11
[SwitchA-pim-vpna] c-rp 11.11.11.11
[SwitchA-pim-vpna] quit
```

Specify VSI-interface 1 as the gateway interface for VSI **vpna**.

```
[SwitchA] vsi vpna
[SwitchA-vsi-vpna] gateway vsi-interface 1
[SwitchA-vsi-vpna] quit
```

4. Configure Switch B:

Enable L2VPN and IP multicast routing.

```

<SwitchB> system-view
[SwitchB] l2vpn enable
[SwitchB] multicast routing
[SwitchB-mrib] quit
# Enable the IGMP snooping feature.
[SwitchB] igmp-snooping
[SwitchB-igmp-snooping] quit
# Disable remote MAC address learning and remote ARP learning.
[SwitchB] vxlan tunnel mac-learning disable
[SwitchB] vxlan tunnel arp-learning disable
# Create VLAN-interface 12 and enter its view.
[SwitchB] vlan 12
[SwitchB-vlan12] quit
[SwitchB] interface vlan-interface 12
# Enable PIM-SM on VLAN-interface 12.
[SwitchB-Vlan-interface12] pim sm
[SwitchB-Vlan-interface12] quit
# Create an EVPN instance on VSI vpnb.
[SwitchB] vsi vpnb
[SwitchB-vsi-vpnb] evpn encapsulation vxlan
[SwitchB-vsi-vpnb-evpn-vxlan] quit
# Enable IGMP snooping and IGMP snooping proxying on VSI vpnb.
[SwitchB-vsi-vpnb] igmp-snooping enable
[SwitchB-vsi-vpnb] igmp-snooping proxy enable
# Create VXLAN 20.
[SwitchB-vsi-vpnb] vxlan 20
[SwitchB-vsi-vpnb-vxlan-20] quit
[SwitchB-vsi-vpnb] quit
# Configure BGP to advertise BGP EVPN routes.
[SwitchB] bgp 200
[SwitchB-bgp-default] peer 4.4.4.4 as-number 200
[SwitchB-bgp-default] peer 4.4.4.4 connect-interface loopback 0
[SwitchB-bgp-default] address-family l2vpn evpn
[SwitchB-bgp-default-evpn] peer 4.4.4.4 enable
[SwitchB-bgp-default-evpn] quit
[SwitchB-bgp-default] quit
# On Ten-GigabitEthernet 1/0/1, create Ethernet service instance 1000 to match VLAN 2.
[SwitchB] interface ten-gigabitethernet 1/0/1
[SwitchB-Ten-GigabitEthernet1/0/1] service-instance 1000
[SwitchB-Ten-GigabitEthernet1/0/1-srv1000] encapsulation s-vid 2
# Map Ethernet service instance 1000 to VSI vpnb.
[SwitchB-Ten-GigabitEthernet1/0/1-srv1000] xconnect vsi vpnb
[SwitchB-Ten-GigabitEthernet1/0/1-srv1000] quit
[SwitchB-Ten-GigabitEthernet1/0/1] quit
# Configure RD and route target settings for VPN instance vpnb.
[SwitchB] ip vpn-instance vpnb
[SwitchB-vpn-instance-vpnb] route-distinguisher 1:1

```

```
[SwitchB-vpn-instance-vpnb] address-family ipv4
[SwitchB-vpn-ipv4-vpnb] vpn-target 2:2
[SwitchB-vpn-ipv4-vpnb] quit
[SwitchB-vpn-instance-vpnb] address-family evpn
[SwitchB-vpn-evpn-vpnb] vpn-target 1:1
[SwitchB-vpn-evpn-vpnb] quit
[SwitchB-vpn-instance-vpnb] quit
```

Configure VSI-interface 1.

```
[SwitchB] interface vsi-interface 1
[SwitchB-Vsi-interfacel] ip binding vpn-instance vpnb
[SwitchB-Vsi-interfacel] ip address 10.1.2.1 255.255.255.0
[SwitchB-Vsi-interfacel] igmp enable
[SwitchB-Vsi-interfacel] mac-address 2-2-2
[SwitchB-Vsi-interfacel] distributed-gateway local
[SwitchB-Vsi-interfacel] local-proxy-arp enable
[SwitchB-Vsi-interfacel] quit
```

Create VSI-interface 2 and configure its L3 VXLAN ID as 1000 to match the traffic from Switch A.

```
[SwitchB] interface vsi-interface 2
[SwitchB-Vsi-interface2] l3-vni 1000
[SwitchB-Vsi-interface2] pim sm
[SwitchB-Vsi-interface2] quit
```

Configure Loopback 1.

```
[SwitchB] interface loopback 1
[SwitchB-LoopBack1] ip binding vpn-instance vpnb
[SwitchB-LoopBack1] ip address 22.22.22.22 32
[SwitchB-LoopBack1] pim sm
[SwitchB-LoopBack1] quit
```

Enable IP multicast routing for VPN instance **vpnb**, and import the traffic of the public instance to VPN instance **vpnb**.

```
[SwitchB] multicast routing vpn-instance vpnb
[SwitchB-mrib-vpnb] multicast extranet select-rpf source 10.1.1.0 24
[SwitchB-mrib-vpnb] multicast extranet select-rpf source 22.22.22.0 24
[SwitchB-mrib-vpnb] quit
```

Create an MDT-based MVXLAN for VPN instance **vpnb** and enter MVXLAN IPv4 address family view. Configure the MVXLAN source interface.

```
[SwitchB] multicast-vpn vxlan vpn-instance vpnb mode mdt
[SwitchB-mvxlan-vpnb] address-family ipv4
[SwitchB-mvxlan-vpnb-ipv4] source loopback 0
[SwitchB-mvxlan-vpnb-ipv4] quit
[SwitchB-mvxlan-vpnb] quit
```

Enter VPN instance PIM view, and configure Loopback 1 as a candidate-BSR and candidate-RP in VPN instance **vpnb**.

```
[SwitchB] pim vpn-instance vpnb
[SwitchB-pim-vpnb] c-bsr 22.22.22.22
[SwitchB-pim-vpnb] c-rp 22.22.22.22
[SwitchB-pim-vpnb] quit
```

Configure Loopback 2.

```

[SwitchB] interface loopback 2
[SwitchB-LoopBack2] ip address 23.23.23.23 32
[SwitchB-LoopBack2] pim sm

# Create IPv4 basic ACL 2000 and enter its view. Create a rule in the ACL to permit only
packets from 225.0.0.0/8.
[SwitchC-acl-ipv4-basic-2000] acl basic 2000
[SwitchC-acl-ipv4-basic-2000] rule permit source 225.0.0.0 0.255.255.255
[SwitchC-acl-ipv4-basic-2000] quit

# Enter PIM view, and configure Loopback 2 as a candidate-RP in the underlay network.
[SwitchC] pim
[SwitchC-pim] c-rp 23.23.23.23 group-policy 2000
[SwitchC-pim] quit

# Specify VSI-interface 1 as the gateway interface for VSI vpnb.
[SwitchB] vsi vpnb
[SwitchB-vsi-vpnb] gateway vsi-interface 1
[SwitchB-vsi-vpnb] quit

```

5. Configure Switch C:

```

# Enable L2VPN and IP multicast routing.
<SwitchC> system-view
[SwitchC] l2vpn enable
[SwitchC] multicast routing
[SwitchC-mrib] quit

# Disable remote MAC address learning and remote ARP learning.
[SwitchC] vxlan tunnel mac-learning disable
[SwitchC] vxlan tunnel arp-learning disable

# Create VLAN-interface 13 and enter its view.
[SwitchC] vlan 13
[SwitchC-vlan13] quit
[SwitchC] interface vlan-interface 13

# Enable PIM-SM on VLAN-interface 13.
[SwitchC-Vlan-interface13] pim sm
[SwitchC-Vlan-interface13] quit

# Create an EVPN instance on VSI vpnc.
[SwitchC] vsi vpnc
[SwitchC-vsi-vpnc] evpn encapsulation vxlan
[SwitchC-vsi-vpnc-evpn-vxlan] quit

# Enable IGMP snooping and IGMP snooping proxying on VSI vpnc.
[SwitchC-vsi-vpnc] igmp-snooping enable
[SwitchC-vsi-vpnc] igmp-snooping proxy enable

# Create VXLAN 30.
[SwitchC-vsi-vpnc] vxlan 30
[SwitchC-vsi-vpnc-vxlan-30] quit
[SwitchC-vsi-vpnc] quit

# Configure BGP to advertise BGP EVPN routes.
[SwitchC] bgp 200
[SwitchC-bgp-default] peer 4.4.4.4 as-number 200
[SwitchC-bgp-default] peer 4.4.4.4 connect-interface loopback 0

```

```
[SwitchC-bgp-default] address-family l2vpn evpn
[SwitchC-bgp-default-evpn] peer 4.4.4.4 enable
[SwitchC-bgp-default-evpn] quit
[SwitchC-bgp-default] quit
```

On Ten-GigabitEthernet 1/0/1, create Ethernet service instance 1000 to match VLAN 3.

```
[SwitchC] interface ten-gigabitethernet 1/0/1
[SwitchC-Ten-GigabitEthernet1/0/1] service-instance 1000
[SwitchC-Ten-GigabitEthernet1/0/1-srv1000] encapsulation s-vid 3
```

Map Ethernet service instance 1000 to VSI vpnc.

```
[SwitchC-Ten-GigabitEthernet1/0/1-srv1000] xconnect vsi vpnc
[SwitchC-Ten-GigabitEthernet1/0/1-srv1000] quit
[SwitchC-Ten-GigabitEthernet1/0/1] quit
```

Configure RD and route target settings for the public instance.

```
[SwitchC] ip public-instance
[SwitchC-public-instance] route-distinguisher 1:1
[SwitchC-public-instance] address-family ipv4
[SwitchC-public-instance-ipv4] vpn-target 2:2
[SwitchC-public-instance-ipv4] quit
[SwitchC-vpn-instance] address-family evpn
[SwitchC-public-instance-evpn] vpn-target 1:1
[SwitchC-public-instance-evpn] quit
[SwitchC-public-instance] quit
```

Create an MDT-based MVXLAN for the public instance and enter MVXLAN IPv4 address family view. Configure the MVXLAN source interface.

```
[SwitchC] multicast-vpn vxlan public-instance mode mdt
[SwitchC-mvxlan-public-instance] address-family ipv4
[SwitchC-mvxlan-public-instance-ipv4] source loopback 0
[SwitchC-mvxlan-public-instance-ipv4] quit
[SwitchC-mvxlan-public-instance] quit
```

Configure VSI-interface 1.

```
[SwitchC] interface vsi-interface 1
[SwitchC-Vsi-interfacel] pim distributed-dr
[SwitchC-Vsi-interfacel] ip address 10.1.3.1 255.255.255.0
[SwitchC-Vsi-interfacel] igmp enable
[SwitchC-Vsi-interfacel] mac-address 3-3-3
[SwitchC-Vsi-interfacel] distributed-gateway local
[SwitchC-Vsi-interfacel] local-proxy-arp enable
[SwitchC-Vsi-interfacel] quit
```

Create VSI-interface 2 and configure its L3 VXLAN ID as 1000 to match the traffic from Switch A.

```
[SwitchC] interface vsi-interface 2
[SwitchC-Vsi-interface3] l3-vni 1000
[SwitchC-Vsi-interface3] pim sm
[SwitchC-Vsi-interface3] quit
```

Configure Loopback 1.

```
[SwitchC] interface loopback 1
[SwitchC-LoopBack1] ip address 33.33.33.33 32
[SwitchC-LoopBack1] pim sm
```

```
[SwitchC-LoopBack1] quit
```

Create IPv4 basic ACL 2000 and enter its view. Create a rule in the ACL to permit only packets from 225.0.0.0/8.

```
[SwitchC-acl-ipv4-basic-2000] acl basic 2000
```

```
[SwitchC-acl-ipv4-basic-2000] rule permit source 225.0.0.0 0.255.255.255
```

```
[SwitchC-acl-ipv4-basic-2000] quit
```

Enter PIM view, and configure Loopback 1 as a candidate-RP in the underlay network.

```
[SwitchC] pim
```

```
[SwitchC-pim] c-rp 33.33.33.33 group-policy 2000
```

```
[SwitchC-pim] quit
```

Specify VSI-interface 1 as the gateway interface for VSI **vpnc.**

```
[SwitchC] vsi vpnc
```

```
[SwitchC-vsi-vpnc] gateway vsi-interface 1
```

```
[SwitchC-vsi-vpnc] quit
```

6. Configure Switch D:

Enable IP multicast routing.

```
<SwitchD> system-view
```

```
[SwitchD] multicast routing
```

```
[SwitchD-mrib] quit
```

Enter PIM view, and configure Loopback 0 as a candidate-BSR and candidate-RP in the underlay network.

```
[SwitchD] pim
```

```
[SwitchD-pim] c-bsr 4.4.4.4
```

```
[SwitchD-pim] c-rp 4.4.4.4
```

```
[SwitchD] quit
```

Enable PIM-SM on VLAN-interface 11.

```
[SwitchD] interface vlan-interface11
```

```
[SwitchD-Vlan-interface11] pim sm
```

```
[SwitchD-Vlan-interface11] quit
```

Enable PIM-SM on VLAN-interface 12.

```
[SwitchD] interface vlan-interface12
```

```
[SwitchD-Vlan-interface12] pim sm
```

```
[SwitchD-Vlan-interface12] quit
```

Enable PIM-SM on VLAN-interface 13.

```
[SwitchD] interface vlan-interface13
```

```
[SwitchD-Vlan-interface13] pim sm
```

```
[SwitchD-Vlan-interface13] quit
```

Establish BGP connections with other transport network switches.

```
[SwitchD] bgp 200
```

```
[SwitchD-bgp-default] group evpn
```

```
[SwitchD-bgp-default] peer 1.1.1.1 group evpn
```

```
[SwitchD-bgp-default] peer 2.2.2.2 group evpn
```

```
[SwitchD-bgp-default] peer 3.3.3.3 group evpn
```

```
[SwitchD-bgp-default] peer evpn as-number 200
```

```
[SwitchD-bgp-default] peer evpn connect-interface loopback 0
```

Configure BGP to advertise BGP EVPN routes, and disable route target filtering of received BGP EVPN routes.

```
[SwitchD-bgp-default] address-family l2vpn evpn
```

```
[SwitchD-bgp-default-evpn] peer evpn enable
[SwitchD-bgp-default-evpn] undo policy vpn-target
# Configure Switch D as an RR.
[SwitchD-bgp-default-evpn] peer evpn reflect-client
[SwitchD-bgp-default-evpn] quit
[SwitchD-bgp-default] quit
```

Verifying the configuration

1. Verify the multicast routing information on Switch A:

Verify that Switch A has multicast routing entries for the public instance.

```
<SwitchA> display pim routing-table
Total 0 (*, G) entries; 2 (S, G) entries

(1.1.1.1, 236.0.0.1)
  RP: 4.4.4.4
  Protocol: pim-sm, Flag: SPT LOC VXLAN_L3
  UpTime: 02:09:52
  Upstream interface: MTunnel0 (VPN: vpna)
    Upstream neighbor: NULL
    RPF prime neighbor: NULL
  Downstream interface information:
  Total number of downstream interfaces: 1
    1: Vlan-interfacell
      Protocol: pim-sm, UpTime: 01:16:34, Expires: 00:03:08

(1.1.1.1, 239.2.0.0)
  RP: 4.4.4.4
  Protocol: pim-sm, Flag: SPT LOC VXLAN_L3
  UpTime: 02:08:52
  Upstream interface: MTunnel1 (VPN: vpna)
    Upstream neighbor: NULL
    RPF prime neighbor: NULL
  Downstream interface information:
  Total number of downstream interfaces: 1
    1: Vlan-interfacell
      Protocol: pim-sm, UpTime: 01:15:34, Expires: 00:03:11
```

Verify that Switch A has multicast routing entries for VPN instance **vpna.**

```
<SwitchA> display pim vpn-instance vpna routing-table
Total 1 (*, G) entries; 1 (S, G) entries

(*, 225.0.0.0)
  RP: 11.11.11.11 (local)
  Protocol: pim-sm, Flag: WC RC
  UpTime: 02:57:31
  Upstream interface: Register-Tunnel0
    Upstream neighbor: NULL
    RPF prime neighbor: NULL
  Downstream interface information:
```

```

Total number of downstream interfaces: 1
  1: MTunnel0
      Protocol: MD, UpTime: 02:57:31, Expires: -

(10.1.1.10, 225.0.0.0)
  RP: 11.11.11.11 (local)
  Protocol: pim-sm, Flag: SPT 2MSDP LOC ACT SQ RC 2MVPN
  UpTime: 04:44:08
  Upstream interface: Vsi-interface1
      Upstream neighbor: NULL
      RPF prime neighbor: NULL
  Downstream interface information:
  Total number of downstream interfaces: 1
    1: MTunnell
        Protocol: MD, UpTime: 02:00:27, Expires: -

```

2. Verify the multicast routing information on Switch B:

Verify that Switch B has multicast routing entries for the public instance.

```
<SwitchB> display pim routing-table
```

```

Total 1 (*, G) entries; 3 (S, G) entries

(*, 225.0.0.0)
  RP: 22.22.22.22 (local)
  Protocol: pim-sm, Flag: WC
  UpTime: 00:01:24
  Upstream interface: Register-Tunnel0
      Upstream neighbor: NULL
      RPF prime neighbor: NULL
  Downstream interface information:
  Total number of downstream interfaces: 1
    1: Extranet (VPN: vpnb)
        Protocol: MD, UpTime: 00:01:32, Expires: -

(10.1.1.10, 225.0.0.0)
  RP: 22.22.22.22 (local)
  Protocol: pim-sm, Flag: SPT ACT 2MVPN FRTNL
  UpTime: 00:00:39
  Upstream interface: Multicast-UPE0 (0.0.0.0)
      Upstream neighbor: NULL
      RPF prime neighbor: NULL
  Downstream interface information:
  Total number of downstream interfaces: 1
    1: Extranet (VPN: vpnb)
        Protocol: MD, UpTime: 00:01:32, Expires: -

(1.1.1.1, 236.2.0.0)
  RP: 4.4.4.4
  Protocol: pim-sm, Flag: SPT
  UpTime: 01:58:46

```

```

Upstream interface: Vlan-interface12
  Upstream neighbor: 12.1.1.4
  RPF prime neighbor: 12.1.1.4
Downstream interface information:
Total number of downstream interfaces: 1
  1: Multicast-UPE0
    Protocol: MD, UpTime: 01:59:16, Expires: -

(1.1.1.1, 239.2.0.0)
  RP: 4.4.4.4
  Protocol: pim-sm, Flag: SPT ACT
  UpTime: 01:58:46
  Upstream interface: Vlan-interface12
    Upstream neighbor: 12.1.1.4
    RPF prime neighbor: 12.1.1.4
  Downstream interface information:
Total number of downstream interfaces: 1
  1: Multicast-UPE0
    Protocol: MD, UpTime: 01:58:46, Expires: -

```

Verify that Switch B has multicast routing entries for VPN instance vpnb.

```

<SwitchB> display pim vpn-instance vpnb routing-table
Total 1 (*, G) entries; 1 (S, G) entries
(*, 225.0.0.0)
  RP: 22.22.22.22 (local)
  Protocol: pim-sm, Flag: WC
  UpTime: 00:01:24
  Upstream interface: Extranet (VPN: public net)
    Upstream neighbor: 127.0.0.1
    RPF prime neighbor: 127.0.0.1
  Downstream interface information:
Total number of downstream interfaces: 1
  1: Vsi-interface1
    Protocol: igmp, UpTime: 00:01:24, Expires: -

(10.1.1.10, 225.0.0.0)
  RP: 22.22.22.22 (local)
  Protocol: pim-sm, Flag: SPT ACT 2MPVN FRTNL
  UpTime: 00:00:39
  Upstream interface: Extranet (VPN: public net)
    Upstream neighbor: 127.0.0.1
    RPF prime neighbor: 127.0.0.1
  Downstream interface information:
Total number of downstream interfaces: 1
  1: Vsi-interface1
    Protocol: pim-sm, UpTime: 00:00:39, Expires: -

```

3. Verify the multicast routing information on Switch C:

Verify that Switch C has multicast routing entries for the public instance.

```

<SwitchC> display pim routing-table

```

```

Total 1 (*, G) entries; 3 (S, G) entries
(*, 225.0.0.0)
  RP: 33.33.33.33 (local)
  Protocol: pim-sm, Flag: WC
  UpTime: 00:01:13
  Upstream interface: Register-Tunnel0
    Upstream neighbor: NULL
    RPF prime neighbor: NULL
  Downstream interface information:
  Total number of downstream interfaces: 1
    1: Vsi-interfacel
      Protocol: igmp, UpTime: 00:01:13, Expires: -

(10.1.1.10, 225.0.0.0)
  RP: 33.33.33.33 (local)
  Protocol: pim-sm, Flag: SPT ACT RQ FRTNL
  UpTime: 00:00:52
  Upstream interface: Multicast-UPE0 (0.0.0.0)
    Upstream neighbor: NULL
    RPF prime neighbor: NULL
  Downstream interface information:
  Total number of downstream interfaces: 1
    1: Vsi-interfacel
      Protocol: pim-sm, UpTime: 00:00:52, Expires: -

(1.1.1.1, 236.2.0.0)
  RP: 4.4.4.4
  Protocol: pim-sm, Flag: SPT
  UpTime: 00:00:30
  Upstream interface: Vlan-interfacel3
    Upstream neighbor: 13.1.1.4
    RPF prime neighbor: 13.1.1.4
  Downstream interface information:
  Total number of downstream interfaces: 1
    1: Multicast-UPE0
      Protocol: MD, UpTime: 01:59:16, Expires: -

(1.1.1.1, 239.2.0.0)
  RP: 4.4.4.4
  Protocol: pim-sm, Flag: SPT ACT
  UpTime: 00:00:30
  Upstream interface: Vlan-interfacel3
    Upstream neighbor: 13.1.1.4
    RPF prime neighbor: 13.1.1.4
  Downstream interface information:
  Total number of downstream interfaces: 1
    1: Multicast-UPE0
      Protocol: MD, UpTime: 01:58:56, Expires: -

```

MVXLAN commands

address-family ipv4

Use **address-family ipv4** to create an MVXLAN IPv4 address family and enter its view, or enter the view of the existing MVXLAN IPv4 address family.

Use **undo address-family ipv4** to delete the MVXLAN IPv4 address family and all settings in MVXLAN IPv4 address family view.

Syntax

```
address-family ipv4
undo address-family ipv4
```

Default

No MVXLAN IPv4 address family exists.

Views

MVXLAN view

Predefined user roles

network-admin

Examples

In MVXLAN view of VPN instance **vpna**, create an MVXLAN IPv4 address family and enter its view.

```
<Sysname> system-view
[Sysname] multicast-vpn vxlan vpn-instance vpna mode mdt
[Sysname-mvxlان-vpna] address-family ipv4
[Sysname-mvxlان-vpna-ipv4]
```

data-delay

Use **data-delay** to set the data-delay period (delay period before the device switches over from the default MDT to the data MDT).

Use **undo data-delay** to restore the default.

Syntax

```
data-delay delay
undo data-delay
```

Default

The data-delay period is 3 seconds.

Views

MVXLAN IPv4 address family view

Predefined user roles

network-admin

Parameters

delay: Specifies a data-delay period in the range of 1 to 60 seconds.

Examples

In MVXLAN IPv4 address family view of VPN instance **vpna**, set the data-delay period to 10 seconds.

```
<Sysname> system-view
[Sysname] multicast-vpn vxlan vpn-instance vpna mode mdt
[Sysname-mvxlan-vpna] address-family ipv4
[Sysname-mvxlan-vpna-ipv4] data-delay 10
```

data-group

Use **data-group** to specify a data group range and optionally configure the criteria for the device to initiate a switchover of the default MDT to a data MDT.

Use **undo data-group** to restore the default.

Syntax

```
data-group group-address { mask-length | mask } [ acl acl-number | threshold threshold-value ] *
```

```
undo data-group
```

Default

No data group range exists, and the device never initiates a switchover of the default MDT to a data MDT.

Views

MVXLAN IPv4 address family view

Predefined user roles

network-admin

Parameters

group-address: Specifies a multicast group address in the range of 224.0.1.0 to 239.255.255.255.

mask-length: Specifies a mask length for the multicast group address. The value range for this argument is 25 to 32.

mask: Specifies a subnet mask for the multicast group address.

acl *acl-number*: Specifies an advanced ACL by its number in the range of 3000 to 3999. If you specify an ACL, the multicast data permitted by the ACL can trigger the switchover. If you do not specify an ACL, any multicast data can trigger the switchover. For the ACL to take effect, specify the protocol type as IP, and include the **source** and **destination** keywords when you create an ACL rule. The **source** and **destination** keywords specify a multicast source address range and a multicast group address range, respectively.

threshold *threshold-value*: Specifies the traffic rate threshold that triggers an MDT switchover. The value range for the *threshold-value* argument is 0 to 16777216 kbps, and the default is 0 kbps.

Usage guidelines

This command enables the device to initiate a switchover of the default MDT to a data MDT in the following situations:

- The multicast data permitted by the ACL keeps arriving for a data-delay period if you specify the **acl** *acl-number* criterion.

- The rate of the multicast data stays higher than the traffic rate threshold during the data-delay period if you specify the **threshold** *threshold-value* criterion.
- If you specify both the **acl** *acl-number* and **threshold** *threshold-value* criteria:
 - The multicast data permitted by the ACL keeps arriving for a data-delay period.
 - The rate of the multicast data stays higher than the traffic rate threshold during the data-delay period.

On a VTEP, the data group range for an MVXLAN cannot include the default group or data groups of any other MVXLAN.

If you execute this command multiple times in the same MVXLAN IPv4 address family view, the most recent configuration takes effect.

Examples

In MVXLAN IPv4 address family view of VPN instance **vpna**, specify 239.1.0.0 through 239.1.0.255 as the data group range.

```
<Sysname> system-view
[Sysname] multicast-vpn vxlan vpn-instance vpna mode mdt
[Sysname-mvxlan-vpna] address-family ipv4
[Sysname-mvxlan-vpna-ipv4] data-group 239.1.0.0 24
```

default-group

Use **default-group** to specify a default group.

Use **undo default-group** to restore the default.

Syntax

```
default-group group-address
undo default-group
```

Default

No default group exists.

Views

MVXLAN IPv4 address family view

Predefined user roles

network-admin

Parameters

group-address: Specifies a default group in the range of 224.0.1.0 to 239.255.255.255.

Usage guidelines

The default group for an MVXLAN must be different from the default group and the data group used by any other MVXLAN.

Examples

In MVXLAN IPv4 address family view of VPN instance **vpna**, specify 236.0.0.1 as the default group.

```
<Sysname> system-view
[Sysname] multicast-vpn vxlan vpn-instance vpna mode mdt
[Sysname-mvxlan-vpna] address-family ipv4
[Sysname-mvxlan-vpna-ipv4] default-group 236.0.0.1
```

display multicast-vpn vxlan data-group receive

Use **display multicast-vpn vxlan data-group receive** to display received data group information in an MVXLAN.

Syntax

```
display multicast-vpn vxlan { vpn-instance instance-name | public-instance } data-group receive [ brief | [ active | group group-address | sender source-address | vpn-source-address [ mask { mask-length | mask } ] | vpn-group-address [ mask { mask-length | mask } ] ]  
* ]
```

Views

Any view

Predefined user roles

network-admin
network-operator

Parameters

vpn-instance *instance-name*: Specifies an MPLS L3VPN instance by its name, a case-sensitive string of 1 to 31 characters.

public-instance: Specifies the public instance.

brief: Displays brief information. If you do not specify this keyword, the command displays detailed information.

active: Specifies data groups that have joined the data MDT.

group *group-address*: Specifies a data group by its IP address in the range of 224.0.1.0 to 239.255.255.255.

sender *source-address*: Specifies an MVXLAN source interface by its IP address.

vpn-source-address: Specifies a multicast source address of the specified VPN instance.

vpn-group-address: Specifies a multicast group address of the specified VPN instance. The value range for this argument is 224.0.1.0 to 239.255.255.255.

mask-length: Specifies a mask length of the specified multicast source address or multicast group address. The value range for this argument is 0 to 32, and the default is 32.

mask: Specifies a subnet mask of the specified multicast source address or multicast group address. The default is 255.255.255.255.

Usage guidelines

If you do not specify any parameters, this command displays detailed information about all received data groups in the MVXLAN of the specified VPN instance or the public instance.

Examples

Display detailed information about received data groups in the MVXLAN of VPN instance **vpna**.

```
<Sysname> display multicast-vpn vxlan vpn-instance vpna data-group receive  
VPN instance: vpna  
Total 2 data-groups for 8 entries  
Total 2 data-groups and 8 entries matched
```

```
Data-group: 226.1.1.0 Reference count: 4 Active count: 2  
Sender: 1.1.1.1 Active count: 1
```

```

(192.6.1.5, 239.1.1.1)      Remaining lifetime: 00:03:10 active
(192.6.1.5, 239.1.1.158)   Remaining lifetime: 00:03:10
Sender: 181.100.1.1    Active count: 1
(195.6.1.2, 239.1.2.12)   Remaining lifetime: 00:03:10 active
(195.6.1.2, 239.1.2.197)  Remaining lifetime: 00:03:10

Data-group: 229.1.1.0    Reference count: 4    Active count: 2
Sender: 185.100.1.1    Active count: 1
(198.6.1.5, 239.1.3.62)   Remaining lifetime: 00:03:10 active
(198.6.1.5, 225.1.1.109)  Remaining lifetime: 00:03:10
Sender: 190.100.1.1    Active count: 1
(200.6.1.2, 225.1.4.80)   Remaining lifetime: 00:03:10 active
(200.6.1.2, 225.1.4.173)  Remaining lifetime: 00:03:10

```

Display brief information about received data groups in the MVXLAN of VPN instance **vpna**.

```

<Sysname> display multicast-vpn vxlan vpn-instance vpna data-group receive brief
VPN instance: vpna
Total 2 data-groups for 8 entries
Total 2 data-groups and 8 entries matched

```

```

Data-group: 226.1.1.0    Reference count: 4    Active count: 2
Data-group: 229.1.1.0    Reference count: 4    Active count: 2

```

Table 2 Command output

| Field | Description |
|---|--|
| Total 2 data-groups for 8 entries | A total of 2 data groups, associated with 8 (S, G) entries. |
| Total 2 data-groups and 8 entries matched | A total of 2 matching data groups, associated with 8 (S, G) entries. |
| Data-group | IP address of the received data group. |
| Reference count | Number of (S, G) entries that use the data group in the VPN instance. |
| Active count | Number of active (S, G) entries (entries with active receivers) that use the data group in the VPN instance. |
| Sender | Address of the VTEP that sent the data group. |
| Remaining lifetime | Remaining lifetime for the (S, G) entry that uses the data group in the VPN instance. |
| active | Indicates that the (S, G) entry is active. |

display multicast-vpn vxlan data-group send

Use **display multicast-vpn vxlan data-group send** to display sent data group information in an MVXLAN.

Syntax

```

display multicast-vpn vxlan { vpn-instance instance-name |
public-instance } data-group send [ group group-address |
vpn-source-address [ mask { mask-length | mask } ] | vpn-group-address
[ mask { mask-length | mask } ] ] *

```

Views

Any view

Predefined user roles

network-admin
network-operator

Parameters

vpn-instance *instance-name*: Specifies an MPLS L3VPN instance by its name, a case-sensitive string of 1 to 31 characters.

public-instance: Specifies the public instance.

group *group-address*: Specifies a data group by its IP address in the range of 224.0.1.0 to 239.255.255.255.

vpn-source-address: Specifies a multicast source address of the specified VPN instance. The value range for this argument is 224.0.1.0 to 239.255.255.255.

vpn-group-address: Specifies a multicast group address of the specified VPN instance. The value range for this argument is 224.0.1.0 to 239.255.255.255.

mask-length: Specifies a mask length of the specified multicast source address or multicast group address. The value range for this argument is 0 to 32, and the default is 32.

mask: Specifies a subnet mask of the specified multicast source address or multicast group address. The default is 255.255.255.255.

Usage guidelines

If you do not specify any parameters, this command displays information about all sent data groups in the MVXLAN of the specified VPN instance or the public instance.

Examples

Display information about all sent data groups in the MVXLAN of VPN instance **vpna**.

```
<Sysname> display multicast-vpn vxlan vpn-instance vpna data-group send
```

```
VPN instance: vpna
```

```
Total 2 data-groups for 6 entries
```

```
Total 2 data-groups and 6 entries matched
```

```
Data-group: 226.1.1.0    Reference count: 3
```

```
(192.6.1.5, 239.1.1.1)    Switchover time: 00:00:21
```

```
(192.6.1.5, 239.1.1.158)  Switchover time: 00:00:21
```

```
(192.6.1.5, 239.1.2.50)   Switchover time: 00:00:05
```

```
Data-group: 226.1.1.1    Reference count: 3
```

```
(192.6.1.2, 225.1.1.1)    Switchover time: 00:00:21
```

```
(192.6.1.2, 225.1.2.50)   Switchover time: 00:00:05
```

```
(192.6.1.5, 239.1.1.159)  Switchover time: 00:00:21
```

Table 3 Command output

| Field | Description |
|---|--|
| Total 2 data-groups for 6 entries | A total of 2 data groups, associated with 6 (S, G) entries. |
| Total 2 data-groups and 6 entries matched | A total of 2 matching data groups, associated with 6 (S, G) entries. |
| Data-group | Data group address. |

| Field | Description |
|-----------------|---|
| Reference count | Number of (S, G) entries that use the data group in the VPN instance. |
| Switchover time | Switchover time of the (S, G) entry that uses the data group in the VPN instance. |

display multicast-vpn vxlan default-group

Use **display multicast-vpn vxlan default-group** to display information about default groups.

Syntax

```
display multicast-vpn vxlan [ vpn-instance instance-name |
public-instance ] default-group { local | remote }
```

Views

Any view

Predefined user roles

network-admin
network-operator

Parameters

vpn-instance *instance-name*: Specifies an MPLS L3VPN instance by its name, a case-sensitive string of 1 to 31 characters.

public-instance: Specifies the public instance.

local: Specifies local default groups.

remote: Specifies remote default groups.

Usage guidelines

If you do not specify the **vpn-instance** *instance-name* option or the **public-instance** keyword, this command displays default group information for all VPN instances and the public instance.

Examples

Display information about local default groups in all VPN instances and the public instance.

```
<Sysname> display multicast-vpn vxlan default-group local
```

MVXLAN local default-group information:

| Group address | Source address | Interface | VPN instance |
|---------------|----------------|-----------|--------------|
| 239.1.1.1 | 1.1.1.1 | MTunnel0 | vpna |
| 239.2.1.1 | 1.1.1.1 | MTunnel1 | vpnb |
| 239.3.1.1 | -- | MTunnel2 | vpnc |

Display information about remote default groups in all VPN instances and the public instance.

```
<Sysname> display multicast-vpn vxlan default-group remote
```

MVXLAN remote default-group information:

| Group address | Source address | Next hop | VPN instance |
|---------------|----------------|----------|--------------|
| 239.1.1.1 | 1.2.0.1 | 1.2.0.1 | a |
| 239.1.1.1 | 1.2.0.2 | 1.2.0.2 | a |
| 239.1.1.1 | 1.2.0.3 | 1.2.0.3 | a |
| 239.1.1.2 | 1.2.0.1 | 1.2.0.1 | b |

| | | | |
|-----------|---------|---------|---|
| 239.1.1.2 | 1.2.0.2 | 1.2.0.2 | b |
| 239.1.1.3 | 1.2.0.1 | 1.2.0.1 | c |

Table 4 Command output

| Field | Description |
|----------------|--|
| Group address | IP address of the default group. |
| Source address | IP address of the MVXLAN source interface. This field displays two hyphens (--) if no source interface IP address is configured for the MVXLAN. |
| Interface | MVXLAN tunnel outgoing interface. |
| Next hop | Next hop address. |
| VPN instance | VPN instance to which the default group belongs. |

multicast extranet select-rpf

Use **multicast extranet select-rpf** to configure an IPv4 MVXLAN extranet RPF selection policy.

Use **undo multicast extranet select-rpf** to delete an IPv4 MVXLAN extranet RPF selection policy.

Syntax

```
multicast extranet select-rpf [ vpn-instance vpn-instance-name ] { group group-address { mask | mask-length } | source source-address { mask | mask-length } } *
```

```
undo multicast extranet select-rpf [ vpn-instance vpn-instance-name ] { group group-address { mask | mask-length } | source source-address { mask | mask-length } } *
```

Default

No IPv4 MVXLAN extranet RPF selection policies are configured.

Views

MRIB view

Predefined user roles

network-admin

Parameters

vpn-instance *vpn-instance-name*: Specifies the MPLS L3VPN instance to which the multicast source belongs. The *vpn-instance-name* argument is a case-sensitive string of 1 to 31 characters. The PIM routing entry will select this VPN instance for the upstream RPF neighbor. If the multicast source belongs to the public network, do not specify this option.

group *group-address*: Specifies a multicast group address in the policy. The address range is 224.0.1.0 to 239.255.255.255. If you do not this option, all multicast traffic that matches the source address in the policy in the source VPN instance will be directed to the receiver VPN instance.

source *source-address*: Specifies a source address in the policy, which can a multicast source address or the address of the RP designated to multicast groups. If you do not specify this option, all multicast traffic that matches the multicast group in the policy in the source VPN instance will be directed to the receiver VPN instance.

mask-length: Specifies a mask length for the multicast source address or the multicast group address. For a multicast group address, the value range for this argument is 4 to 32. For a multicast source address, the value range for this argument is 0 to 32.

mask: Specifies a subnet mask for the multicast source address or the multicast group address.

Usage guidelines

MVXLAN extranet RPF routing policies are used for multicast transmission when multicast sources and receivers are located in different VPNs.

The PIM mode in the source VPN instance and the receiver VPN instance must be the same. Only PIM-SM and PIM-SSM are supported.

As a best practice, if PIM-SM is used for MVXLAN extranet solution to implement intra-VPN multicast transmission for multiple multicast groups, configure a dedicated RP for these groups. In addition, specify the address of the RP as the source address in RPF selection policies.

In PIM-SM mode, you can configure only one RPF selection policy for a multicast source and group in a VPN instance.

To implement source-specific RPF selection in MVXLAN extranet, you must configure two MVXLAN extranet RPF routing policies as follows:

- In one policy, specify the address of the RP designated to the multicast group that requires inter-VPN multicast communication as the source address.
- In the other policy, specify the multicast source in the source VPN instance as the source address.

To implement source-and-group-specific RPF selection in MVXLAN extranet, you must configure two MVXLAN extranet RPF routing policies as follows:

- In one policy, specify the address of the RP designated to the multicast group as the source address, and specify the multicast group.
- In the other policy, specify the multicast source in the source VPN instance as the source address, and specify the multicast group.

If an MVXLAN extranet RPF selection policy with only the multicast group address specified is configured in the receiver VPN instance, the intra-VPN multicast traffic transmission will be interrupted.

Multicast source addresses in different MVXLAN extranet RPF routing policies cannot be the same, but they can overlap. The same restriction applies to the multicast group addresses in different MVXLAN extranet RPF routing policies. If multiple routing policies exist for an (S, G) entry, the device selects the policy in which the multicast group address has the longest mask. If multiple policies have the same mask length, the device selects the policy in which the multicast source address has the longest mask.

Examples

Configure an IPv4 MVXLAN extranet RPF selection policy to direct multicast traffic sourced from 1.1.1.1 from VPN instance **vpn-source** to VPN instance **vpn-receiver**.

```
<Sysname> system-view
[Sysname] multicast routing vpn-instance vpn-receiver
[Sysname-mrib-vpn-receiver] multicast extranet select-rpf vpn-instance vpn-source source
192.168.10.0 24
[Sysname-mrib-vpn-receiver] multicast extranet select-rpf vpn-instance vpn-source source
1.1.1.1 24
```

multicast-vpn vxlan public-instance

Use **multicast-vpn vxlan public-instance** to create an MVXLAN for the public instance and enter its view, or enter the view of the existing public instance MVXLAN.

Use **undo multicast-vpn vxlan public-instance** to delete the public instance MVXLAN and all settings in public instance MVXLAN view.

Syntax

```
multicast-vpn vxlan public-instance mode mdt
undo multicast-vpn vxlan public-instance
```

Default

No MVXLAN exists in the public instance.

Views

System view

Predefined user roles

network-admin

Parameters

mode: Specifies the MVXLAN mode.

mdt: Sets the mode to MDT.

Examples

Create an MDT-based MVXLAN for the public instance and enter public instance MVXLAN view.

```
<Sysname> system-view
```

```
[Sysname] multicast-vpn vxlan public-instance mode mdt
```

```
[Sysname-mvxlan-public-instance]
```

multicast-vpn vxlan vpn-instance

Use **multicast-vpn vxlan vpn-instance** to create an MVXLAN for a VPN instance and enter its view, or enter the view of an existing VPN instance MVXLAN.

Use **undo multicast-vpn vxlan vpn-instance** to delete a VPN instance MVXLAN and all settings in its view.

Syntax

```
multicast-vpn vxlan vpn-instance instance-name mode mdt
undo multicast-vpn vxlan vpn-instance instance-name
```

Default

No MVXLAN exists in a VPN instance.

Views

System view

Predefined user roles

network-admin

Parameters

instance-name: Specifies an MPLS L3VPN instance by its name, a case-sensitive string of 1 to 31 characters.

mode: Specifies the MVXLAN mode.

mdt: Sets the mode to MDT.

Examples

Create an MDT-based MVXLAN for VPN instance **vpna** and enter VPN instance MVXLAN view.

```
<Sysname> system-view
[Sysname] multicast-vpn vxlan vpn-instance vpna mode mdt
[Sysname-mvxlan-vpna]
```

pim distributed-dr

Use **pim distributed-dr** to configure a VSI interface as a distributed DR interface.

Use **undo pim distributed-dr** to restore the default.

Syntax

```
pim distributed-dr
undo pim distributed-dr
```

Default

A VSI instance is not a distributed DR interface.

Views

VSI interface view

Predefined user roles

network-admin

Usage guidelines

On the VTEPs configured with MVXLAN, you must specify the VSI interfaces that act as distributed EVPN gateways as distributed DR interfaces. This operation ensures that a distributed EVPN gateway can forward multicast traffic to the local site.

Examples

Configure VSI-interface 1 as a distributed DR interface.

```
<Sysname> system-view
[Sysname] interface vsi-interface 1
[Sysname-Vsi-interfacel] pim distributed-dr
```

source

Use **source** to specify an MVXLAN source interface.

Use **undo source** to restore the default.

Syntax

```
source interface-type interface-number
undo source
```

Default

No MVXLAN source interface is specified.

Views

MVXLAN IPv4 address family view

Predefined user roles

network-admin

Parameters

interface-type interface-number: Specify an interface by its type and number.

Usage guidelines

For the VTEP to obtain correct routing information, you must specify the interface used for establishing BGP peer relationship as the MVXLAN source interface.

Examples

```
# In MVXLAN IPv4 address family view of VPN instance vpna, specify Loopback 0 as the MVXLAN source interface. (Loopback 0 is the source interface used for establishing BGP peer relationship.)
<Sysname> system-view
[Sysname] multicast-vpn vxlan vpn-instance vpna mode mdt
[Sysname-mvxlan-vpna] address-family ipv4
[Sysname-mvxlan-vpna-ipv4] source loopback 0
```

Modified feature: Configuring LLDP to advertise IPv4 and IPv6 management addresses at the same time

Feature change description

From this software version, the device supports advertising the IPv4 and IPv6 addresses of the LLDP frame sending ports as the management addresses at the same time. If you specify the IPv6 address, the management address is the IPv6 address. If you do not specify the IPv6 address, the management address is the IPv4 address. If you configure the IP addresses of the same type multiple times, the most recent configuration takes effect.

Command changes

Modified command: `lldp global tlv-enable basic-tlv management-address-tlv`

Old syntax

```
lldp [ agent { nearest-customer | nearest-nontpmr } ] global tlv-enable
basic-tlv management-address-tlv [ ipv6 ] { ip-address | interface loopback
interface-number | interface m-gigabitethernet interface-number |
interface vlan-interface interface-number }

undo lldp [ agent { nearest-customer | nearest-nontpmr } ] global tlv-enable
basic-tlv management-address-tlv
```

New syntax

```
lldp [ agent { nearest-customer | nearest-nontpmr } ] global tlv-enable
basic-tlv management-address-tlv [ ipv6 ] { ip-address | interface loopback
interface-number | interface m-gigabitethernet interface-number |
interface vlan-interface interface-number }

undo lldp [ agent { nearest-customer | nearest-nontpmr } ] global tlv-enable
basic-tlv management-address-tlv [ [ ipv6 ] { ip-address | interface
loopback interface-number | interface m-gigabitethernet interface-number
| interface vlan-interface interface-number } ]
```

Views

System view

Change description

Before modification: The IPv4 and IPv6 management addresses cannot be advertised at the same time. After you enable advertisement of the management address TLV globally, only the management address that is configured most recently can be advertised. If you execute this command multiple times, the most recent configuration takes effect.

After modification: The IPv4 and IPv6 management addresses can be advertised at the same time. If you configure the IP addresses of the same type multiple times, the most recent configuration takes effect. When you disable advertisement of the management address TLV globally, you can disable IPv4 and IPv6 management addresses separately.

Modified feature: Specifying a padding mode for the Interface-ID option on the DHCPv6 relay agent

Feature change description

This release added support for user MAC padding mode for the Interface-ID option. The DHCPv6 relay agent in this mode pads the Interface-ID option with the MAC address of the DHCPv6 client.

Command changes

Modified command: `ipv6 dhcp relay interface-id`

Old syntax

```
ipv6 dhcp relay interface-id { bas | interface }
```

New syntax

```
ipv6 dhcp relay interface-id { bas | interface | user-mac }
```

Views

Interface view

Change description

Before modification: The command does not support the **user-mac** keyword.

After modification: The **user-mac** keyword was added to the command. The **user-mac** keyword specifies the user MAC mode. The relay agent in this mode pads the Interface-ID option with the MAC address of the DHCPv6 client.

Modified feature: Displaying monitor link group information

Feature change description

From this release, the **DR DOWN (Monitor Link)** value, which means that a DR interface is shut down by Monitor Link, is added to the **Status** field for member interfaces.

Command changes

Modified command: display monitor-link group

Syntax

```
display monitor-link group { group-id | all }
```

Views

Any view

Change description

Before modification: The **DR DOWN (Monitor Link)** value is not supported for the **Status** field.

After modification: The **DR DOWN (Monitor Link)** value is supported for the **Status** field.

Modified feature: Displaying the path that the IPv6 packets traverse from source to destination

Feature change description

From this release, you can use the **tracert ipv6** command to display the path that the IPv6 packets traverse from a specified source IPv6 address to a destination. The source IPv6 address must be a valid IPv6 address on the device.

Command changes

Modified command: tracert ipv6

Old syntax

```
tracert ipv6 [ -f first-hop | -m max-hops | -p port | -q packet-number | -t traffic-class | -vpn-instance vpn-instance-name [ -resolve-as { global | none | vpn } ] | -w timeout ] * host
```

New syntax

```
tracert ipv6 [ -a source-ipv6 | -f first-hop | -m max-hops | -p port | -q packet-number | -t traffic-class | -vpn-instance vpn-instance-name [ -resolve-as { global | none | vpn } ] | -w timeout ] * host
```

Views

Any view

Change description

The **-a *source-ipv6*** option is added to this command. It specifies a source IPv6 address for the IPv6 probe packets. The address must be a valid IPv6 address on the device. If you do not specify a source IPv6 address, the source IPv6 address of the IPv6 probe packets is the IPv6 address of the outbound interface.

Modified feature: Exchanging both IPv4 and IPv6 routes with a BGP peer

Feature change description

From this release, BGP can exchange IPv6 routes with an IPv4 peer and exchange IPv4 routes with an IPv6 peer. You can specify IPv6 peers for related commands in IPv4 address family view and specify IPv4 peers for related commands in IPv6 address family view.

Command changes

Syntax

```
display bgp peer
display bgp peer received prefix-list
display bgp routing-table ipv4 unicast
display bgp update-group
refresh bgp
reset bgp
peer additional-paths
peer advertise additional-paths best
peer advertise origin-as-validation
peer advertise-community
peer advertise-ext-community
peer advertise-policy exist-policy
peer advertise-policy non-exist-policy
peer allow-as-loop
peer as-path-acl
peer capability-advertise orf prefix-list
peer default-route-advertise
peer enable
peer filter-policy
peer keep-all-routes
peer next-hop-local
peer preferred-value
peer prefix-list
peer public-as-only
peer reflect-client
peer route-limit
peer route-policy
peer soo
```

Change description

Before modification:

- BGP cannot exchange IPv6 routes with IPv4 peers.
- When you execute the following commands, you cannot specify an IPv6 peer in BGP IPv4 unicast address family view or BGP-VPN IPv4 unicast address family view, or specify an IPv4 peer in BGP-VPN IPv6 unicast address family view
 - `peer additional-paths`
 - `peer advertise additional-paths best`
 - `peer advertise origin-as-validation`
 - `peer advertise-community`
 - `peer advertise-ext-community`
 - `peer allow-as-loop`
 - `peer as-path-acl`
 - `peer capability-advertise orf prefix-list`
 - `peer default-route-advertise`
 - `peer enable`
 - `peer filter-policy`
 - `peer keep-all-routes`
 - `peer next-hop-local`
 - `peer preferred-value`
 - `peer prefix-list`
 - `peer public-as-only`
 - `peer reflect-client`
 - `peer route-limit`
 - `peer route-policy`
 - `peer soo`
- You cannot specify an IPv6 peer when you execute the following commands in BGP IPv4 unicast address family view:
 - `peer advertise-policy exist-policy`
 - `peer advertise-policy non-exist-policy`
- The following commands do not support specifying IPv6 peers when displaying information for IPv4 address family or IPv4 address family in a VPN instance or specifying IPv4 peers when displaying information for IPv6 address family in a VPN instance:
 - `display bgp peer`
 - `display bgp peer received prefix-list`
 - `display bgp routing-table ipv4 unicast`
 - `display bgp update-group`
- The following commands do not support specifying IPv6 peers when resetting information for IPv4 address family or IPv4 address family in a VPN instance or specifying IPv4 peers when resetting information for IPv6 address family in a VPN instance:
 - `refresh bgp`
 - `reset bgp`

After modification:

- BGP can exchange IPv6 routes with IPv4 peers.

- When you execute the following commands, you can specify an IPv6 peer in BGP IPv4 unicast address family view or BGP-VPN IPv4 unicast address family view, and specify an IPv4 peer in BGP-VPN IPv6 unicast address family view
 - `peer additional-paths`
 - `peer advertise additional-paths best`
 - `peer advertise origin-as-validation`
 - `peer advertise-community`
 - `peer advertise-ext-community`
 - `peer allow-as-loop`
 - `peer as-path-acl`
 - `peer capability-advertise orf prefix-list`
 - `peer default-route-advertise`
 - `peer enable`
 - `peer filter-policy`
 - `peer keep-all-routes`
 - `peer next-hop-local`
 - `peer preferred-value`
 - `peer prefix-list`
 - `peer public-as-only`
 - `peer reflect-client`
 - `peer route-limit`
 - `peer route-policy`
 - `peer soo`
- You can specify an IPv6 peer when you execute the following commands in BGP IPv4 unicast address family view:
 - `peer advertise-policy exist-policy`
 - `peer advertise-policy non-exist-policy`
- The following commands support specifying IPv6 peers when displaying information for IPv4 address family or IPv4 address family in a VPN instance and specifying IPv4 peers when displaying information for IPv6 address family in a VPN instance:
 - `display bgp peer`
 - `display bgp peer received prefix-list`
 - `display bgp routing-table ipv4 unicast`
 - `display bgp update-group`
- The following commands support specifying IPv6 peers when resetting information for IPv4 address family or IPv4 address family in a VPN instance and specifying IPv4 peers when resetting information for IPv6 address family in a VPN instance:
 - `refresh bgp`
 - `reset bgp`

Modified feature: Enabling load balancing for BGP routes that have different AS_PATH attributes of the same length

Feature change description

From this release, when enabling load balancing for routes that have different AS_PATH attributes of the same length, you can specify the following keywords:

- **ebgp**: Enables load balancing for EBGp routes that have different AS_PATH attributes of the same length.
- **ibgp**: Enables load balancing for IBGP routes that have different AS_PATH attributes of the same length.

If you do not specify the **ibgp** or **ebgp** keyword, the following rules apply:

- BGP performs load balancing for EBGp routes that have different AS_PATH attributes of the same length.
- BGP performs load balancing for IBGP routes that have different AS_PATH attributes of the same length.
- BGP does not perform load balancing between EBGp and IBGP routes that have different AS_PATH attributes of the same length.

Command changes

Modified command: balance as-path-relax

Use **balance as-path-relax** to enable load balancing for routes that have different AS_PATH attributes of the same length.

Use **undo balance as-path-relax** to restore the default.

Old syntax

```
balance as-path-relax
undo balance as-path-relax
```

New syntax

```
balance as-path-relax [ ebgp | ibgp ]
undo balance as-path-relax [ ebgp | ibgp ]
```

Parameters

ebgp: Enables load balancing for EBGp routes that have different AS_PATH attributes of the same length.

ibgp: Enables load balancing for IBGP routes that have different AS_PATH attributes of the same length.

Change description

Before modification: The command does not support the **ebgp** or **ibgp** keyword.

After modification: The command supports the **ebgp** and **ibgp** keywords.

Modified feature: Enabling BGP to exchange labeled routes with a peer or peer group

Feature change description

From this release, the device supports enabling labeled route exchange with a peer in BGP-VPN IPv6 unicast address family view.

Command changes

Modified command: peer label-route-capability

Syntax

In BGP IPv4 unicast address family view/BGP-VPN IPv4 unicast address family view/BGP IPv6 unicast address family view:

```
peer { group-name | ipv4-address [ mask-length ] } label-route-capability  
undo peer { group-name | ipv4-address [ mask-length ] }  
label-route-capability
```

In BGP-VPN IPv6 unicast address family view:

```
peer ipv4-address [ mask-length ] label-route-capability  
undo peer ipv4-address [ mask-length ] label-route-capability
```

Views

BGP IPv4 unicast address family view

BGP-VPN IPv4 unicast address family view

BGP IPv6 unicast address family view

BGP-VPN IPv6 unicast address family view

Change description

Before modification: The device does not support executing this command in BGP-VPN IPv6 unicast address family view.

After modification: The device supports executing this command in BGP-VPN IPv6 unicast address family view.

Modified feature: Configuring a static neighbor entry

Feature change description

This release added support for configuring static entries for neighbors that are connected to the device through VSI interfaces.

Command changes

Modified command: ipv6 neighbor

Use **ipv6 neighbor** to configure a static neighbor entry.

Use **undo ipv6 neighbor** to delete a neighbor entry.

Old command

```
ipv6 neighbor ipv6-address mac-address { vlan-id port-type port-number |  
interface interface-type interface-number } [ vpn-instance  
vpn-instance-name ]
```

```
undo ipv6 neighbor ipv6-address interface-type interface-number
```

New command

```
ipv6 neighbor ipv6-address mac-address { vlan-id port-type port-number |  
interface interface-type interface-number | vsi-interface  
vsi-interface-id tunnel number vsi vsi-name | vsi-interface  
vsi-interface-id interface-type interface-number service-instance  
instance-id vsi vsi-name } [ vpn-instance vpn-instance-name ]
```

```
undo ipv6 neighbor ipv6-address interface-type interface-number
```

Default

No static neighbor entries exist.

Views

System view

Parameters

ipv6-address: Specifies the IPv6 address of the static neighbor entry.

mac-address: Specifies the MAC address (48 bits) of the static neighbor entry, in the format of H-H-H.

vlan-id: Specifies the VLAN ID of the static neighbor entry, in the range of 1 to 4094.

port-type port-number: Specifies a Layer 2 port of the static neighbor entry by its type and number.

interface *interface-type interface-number*: Specifies a Layer 3 interface of the static neighbor entry by its type and number.

vsi-interface *vsi-interface-id*: Specifies an input VSI interface for packets received from the neighbor in the entry. The *vsi-interface-id* argument specifies the VSI interface number.

tunnel *number*: Specifies an output tunnel interface for packets sent to the neighbor in the entry. The *number* argument specifies the tunnel interface number.

interface-type interface-number: Specifies a Layer 2 interface by its type and number. The device determines an output interface for packets sent to the neighbor in the entry based on the specified Layer 2 interface and Ethernet service instance.

vsi *vsi-name*: Specifies a VSI name, a case-sensitive string of 1 to 31 characters.

service-instance *instance-id*: Specifies the Ethernet service instance of the entry. The *instance-id* specifies the Ethernet service instance ID in the range of 1 to 4096. You must specify this option if a Layer 2 interface is specified. This option is not configurable if an interface of another type is specified.

vpn-instance *vpn-instance-name*: Specifies an MPLS L3VPN instance to which the static neighbor entry belongs. The *vpn-instance-name* argument represents the VPN instance name, a case-sensitive string of 1 to 31 characters. If you do not specify a VPN instance, this command configures a static neighbor entry for the public network.

Change description

Before modification: The device does not support configuring static entries for neighbors that are connected to the device through VSI interfaces.

After modification: The device supports configuring static entries for neighbors that are connected to the device through VSI interfaces.

If the device and its neighbor are connected through a VSI interface, use one of the following methods to configure the neighbor entry:

- Method 1—Specify a neighbor IPv6 address and MAC address, and specify the input interface (VSI interface), output interface (tunnel interface), and VSI name by using the `vsi-interface vsi-interface-id tunnel number vsi vsi-name` option.
- Method 2—Specify a neighbor IPv6 address, MAC address, input interface (VSI interface), output interface (determined by a Layer 2 interface and Ethernet service instance), and VSI name. Specify the input interface, output interface, and VSI name by using the `vsi-interface vsi-interface-id interface-type interface-number service-instance instance-id vsi vsi-name` option.

Modified command: display ipv6 neighbors

Use `display ipv6 neighbors` to display IPv6 neighbor information.

Syntax

```
display ipv6 neighbors { { ipv6-address | all | dynamic | static } [ slot slot-number ] | interface interface-type interface-number | vlan vlan-id }  
[ verbose ]
```

Views

Any view

Change description

Before modification: The output from the `display ipv6 neighbors all verbose` command does not include the **VXLAN ID**, **VSI name**, and **VSI interface** fields.

After modification: The **VXLAN ID**, **VSI name**, and **VSI interface** fields were added to the output from the `display ipv6 neighbors all verbose` command.

Modified feature: Modifying information in received BGP EVPN IP prefix advertisement routes on an ED

Feature change description

As from this version, an ED supports replacing the route targets of received BGP EVPN IP prefix advertisement routes with those of the matching local VPN instance.

Command changes

Modified command: peer re-originated

Old syntax

```
peer { group-name | ipv4-address [ mask-length ] | ipv6-address  
[ prefix-length ] } re-originated
```

```
undo peer { group-name | ipv4-address [ mask-length ] | ipv6-address  
[ prefix-length ] } re-originated
```

New syntax

```
peer { group-name | ipv4-address [ mask-length ] | ipv6-address  
[ prefix-length ] } re-originated [ replace-rt ]  
undo peer { group-name | ipv4-address [ mask-length ] | ipv6-address  
[ prefix-length ] } re-originated
```

Views

BGP EVPN address family view

Change description

The **replace-rt** keyword was added to this command. This keyword replaces the L3 VXLAN ID, RD, and route targets of IP prefix advertisement routes with those of the matching local VPN instance. If you do not specify this keyword, only the L3 VXLAN ID and RD will be replaced.

Modified feature: Configuring passive TCP connection settings for a VXLAN OVSDb VTEP

Feature change description

As from this version, the OVSDb server on a VXLAN OVSDb VTEP supports using a basic IPv4 ACL to filter OVSDb TCP connection requests.

Command changes

Modified command: ovbdb server ptcp

Old syntax

```
ovbdb server ptcp [ port port-number ]
```

New syntax

```
ovbdb server ptcp [ port port-number ] [ acl acl-number ]
```

Views

System view

Change description

The **acl acl-number** option was added to specify a basic IPv4 ACL by its ACL number to filter OVSDb TCP connection requests. The value range for the *acl-number* argument is 2000 to 2999.

If you specify a basic IPv4 ACL for this command, the OVSDb server accepts only the OVSDb TCP connection requests that are permitted by the ACL. The OVSDb server rejects any OVSDb TCP connection requests that do not match the ACL or that are denied by the ACL.

The OVSDb server does not filter OVSDb TCP connection requests in any of the following situations:

- No IPv4 ACL is specified for this command.
- The specified IPv4 ACL does not exist.
- The specified IPv4 ACL does not have rules.

Modification to rules in the specified IPv4 ACL takes effect only on OVSDB TCP connection requests that are received after the modification. If you specify another IPv4 ACL, the IPv4 ACL also takes effect only on OVSDB TCP connection requests that are received after the modification.

Modified feature: Setting the password for local password authentication

Feature change description

The minimum length of a password in plaintext form was changed from one character to four characters, and the password must contain a minimum of two character types.

Command changes

Modified command: set authentication password

Syntax

```
set authentication password { hash | simple } string  
undo set authentication password
```

Views

User line view

User line class view

Change description

Before modification: The password in plaintext form is a string of 1 to 16 characters.

After modification: The password in plaintext form is a string of 4 to 16 characters, and must contain a minimum of two character types.

Modified feature: Enabling an ED to replace the L3 VXLAN ID, RD, and route targets of BGP EVPN routes

Enabling an ED to replace the L3 VXLAN ID, RD, and route targets of BGP EVPN routes

About this task

In an EVPN-DCI network, use this feature to hide the L3 VXLAN IDs of data centers or enable communication between data centers that use different L3 VXLAN IDs or route targets.

After you enable this feature on an ED, the ED performs the following operations after receiving BGP EVPN routes:

1. Matches the route targets of the routes with the import route targets of local VPN instances.
2. Replaces the L3 VXLAN ID, RD, and route targets of the routes with those of the matching local VPN instance.
3. Advertises the routes to a VTEP or remote ED.

After you execute the **peer re-originated** command, the ED advertises only reoriginated BGP EVPN routes. For the ED to advertise both original and reoriginated BGP EVPN routes, execute the **peer advertise original-route** command.

An ED configured with the **peer re-originated** and **peer advertise original-route** commands advertises both original and reoriginated BGP EVPN routes. For the ED to advertise only original BGP EVPN routes, execute the **peer suppress re-originated** command on the ED.

Restrictions and guidelines

If the RD of a received BGP EVPN route is identical to the RD of the matching local VPN instance, an ED does not replace the L3 VXLAN ID and route targets of the route or reoriginate the route. As a result, the ED does not advertise the route. As a best practice, assign unique RDs to VPN instances on different EVPN gateways and EDs when you use this feature.

Procedure

1. Enter system view.
system-view
2. Enter BGP instance view.
bgp as-number [instance instance-name]
3. Enter BGP EVPN address family view.
address-family l2vpn evpn
4. Replace the L3 VXLAN ID, RD, and route targets (optional) of received BGP EVPN routes.
peer { group-name | ipv4-address [mask-length] | ipv6-address [prefix-length] } re-originated [imet | ip-prefix | mac-ip] [replace-rt]

By default, the device does not modify the BGP EVPN routes that are received from peers or peer groups.
5. (Optional.) Enable the device to advertise original BGP EVPN routes together with the reoriginated BGP EVPN routes after the **peer re-originated** command is executed.
peer { group-name | ipv4-address [mask-length] | ipv6-address [prefix-length] } advertise original-route

By default, the device advertises only reoriginated BGP EVPN routes to peers and peer groups after the **peer re-originated** command is executed.
6. (Optional.) Suppress advertisement of reoriginated BGP EVPN routes to a peer or peer group.
peer { group-name | ipv4-address [mask-length] | ipv6-address [prefix-length] } suppress re-originated { imet | ip-prefix | mac-ip }

By default, the device advertises reoriginated BGP EVPN routes to peers and peer groups after the **peer re-originated** command is executed.

Command reference

Modified command: peer re-originated

Old syntax

```
peer { group-name | ipv4-address [ mask-length ] | ipv6-address [ prefix-length ] } re-originated [ replace-rt ]  
  
undo peer { group-name | ipv4-address [ mask-length ] | ipv6-address [ prefix-length ] } re-originated
```

New syntax

```
peer { group-name | ipv4-address [ mask-length ] | ipv6-address [ prefix-length ] } re-originated [ imet | ip-prefix | mac-ip ] [ replace-rt ]
```

```
undo peer { group-name | ipv4-address [ mask-length ] | ipv6-address [ prefix-length ] } re-originated [ imet | ip-prefix | mac-ip ]
```

Views

BGP EVPN address family view

Change description

The **imet**, **ip-prefix**, and **mac-ip** keywords were added to this command.

imet: Specifies inclusive multicast Ethernet tag (IMET) routes.

ip-prefix: Specifies IP prefix advertisement routes.

mac-ip: Specifies MAC/IP advertisement routes.

New command: peer advertise original-route

Use **peer advertise original-route** to enable the device to advertise original BGP EVPN routes to a peer or peer group.

Use **undo peer advertise original-route** to disable the device from advertising original BGP EVPN routes to a peer or peer group.

Syntax

```
peer { group-name | ipv4-address [ mask-length ] | ipv6-address [ prefix-length ] } advertise original-route
```

```
undo peer { group-name | ipv4-address [ mask-length ] | ipv6-address [ prefix-length ] } advertise original-route
```

Default

The device advertises only reoriginated BGP EVPN routes to peers and peer groups after the **peer re-originated** command is executed.

Views

BGP EVPN address family view

Predefined user roles

network-admin

Parameters

group-name: Specifies a peer group by its name, a case-sensitive string of 1 to 47 characters. The peer group must exist.

ipv4-address: Specifies a peer by its IPv4 address. The peer must exist.

mask-length: Specifies a mask length in the range of 0 to 32. To specify a subnet, you must specify both the *ipv4-address* and *mask-length* arguments.

ipv6-address: Specifies a peer by its IPv6 address. The peer must exist.

prefix-length: Specifies a prefix length in the range of 0 to 128. To specify a subnet, you must specify both the *ipv6-address* and *prefix-length* arguments.

Usage guidelines

For this command to take effect on an ED, first execute the **peer re-originated** command.

In an EVPN-DCI network, an ED configured with the **peer re-originated** command advertises only reoriginated BGP EVPN routes. For the ED to advertise both original and reoriginated BGP EVPN routes to a peer or peer group, execute the **peer advertise original-route** command on the ED.

Examples

Enable the device to advertise original BGP EVPN routes to peer 1.1.1.1.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] address-family l2vpn evpn
[Sysname-bgp-default-evpn] peer 1.1.1.1 advertise original-route
```

Related commands

peer re-originated

peer suppress re-originated

New command: peer suppress re-originated

Use **peer suppress re-originated** to suppress advertisement of reoriginated BGP EVPN routes to a peer or peer group.

Use **undo peer suppress re-originated** to disable suppression of reoriginated BGP EVPN route advertisement to a peer or peer group.

Syntax

```
peer { group-name | ipv4-address [ mask-length ] | ipv6-address [ prefix-length ] } suppress re-originated { imet | ip-prefix | mac-ip }
undo peer { group-name | ipv4-address [ mask-length ] | ipv6-address [ prefix-length ] } suppress re-originated { imet | ip-prefix | mac-ip }
```

Default

The device advertises reoriginated BGP EVPN routes to peers and peer groups after the **peer re-originated** command is executed.

Views

BGP EVPN address family view

Predefined user roles

network-admin

Parameters

group-name: Specifies a peer group by its name, a case-sensitive string of 1 to 47 characters. The peer group must exist.

ipv4-address: Specifies a peer by its IPv4 address. The peer must exist.

mask-length: Specifies a mask length in the range of 0 to 32. To specify a subnet, you must specify both the *ipv4-address* and *mask-length* arguments.

ipv6-address: Specifies a peer by its IPv6 address. The peer must exist.

prefix-length: Specifies a prefix length in the range of 0 to 128. To specify a subnet, you must specify both the *ipv6-address* and *prefix-length* arguments.

imet: Specifies inclusive multicast Ethernet tag (IMET) routes.

ip-prefix: Specifies IP prefix advertisement routes.

mac-ip: Specifies MAC/IP advertisement routes.

Usage guidelines

An ED configured with the **peer re-originated** and **peer advertise original-route** commands advertises both original and reoriginated BGP EVPN routes. For the ED to advertise only original BGP EVPN routes to a peer or peer group, execute the **peer suppress re-originated** command on the ED.

Examples

```
# Suppress advertisement of reoriginated IP prefix advertisement routes to peer 1.1.1.1.
```

```
<Sysname> system-view
```

```
[Sysname] bgp 100
```

```
[Sysname-bgp-default] address-family l2vpn evpn
```

```
[Sysname-bgp-default-evpn] peer 1.1.1.1 suppress re-originated ip-prefix
```

Related commands

peer advertise original-route

peer re-originated

Modified feature: Displaying Ethernet service instance information

Feature change description

In this version, the value for the **Type** field was changed from **Dynamic (MLAG)** to **Dynamic (DRNI)** for dynamic Ethernet service instances created on the IPL in an EVPN DR system.

Command changes

Modified command: display l2vpn service-instance

Syntax

```
display l2vpn service-instance [ interface interface-type  
interface-number service-instance instance-id ] [ verbose ]
```

Views

Any view

Change description

Before modification: When you display detailed information about Ethernet service instances in an EVPN DR system, the value for the **Type** field is **Dynamic (MLAG)** for dynamic Ethernet service instances created on the IPL.

After modification: When you display detailed information about Ethernet service instances in an EVPN DR system, the value for the **Type** field is **Dynamic (DRNI)** for dynamic Ethernet service instances created on the IPL.

Modified feature: Configuring OSPF FRR

Feature change description

This release added support for calculating a backup next hop through LFA calculation for ECMP routes.

Command changes

Modified command: fast-reroute

Old syntax

```
fast-reroute { lfa [ abr-only ] | route-policy route-policy-name }  
undo fast-reroute
```

New syntax

```
fast-reroute { lfa [ abr-only | ecmp-shared ] | route-policy  
route-policy-name }  
undo fast-reroute
```

Change description

Before modification: The command does not support the **ecmp-shared** keyword.

After modification: The **ecmp-shared** keyword was added to the command. After you specify the **ecmp-shared** keyword for the command, OSPF calculates a backup next hop through LFA calculation for all routes, including ECMP routes. ECMP routes share one backup next hop. If the next hops of all ECMP routes fail, OSPF uses the shared backup next hop to forward packets.

Modified feature: Displaying routing table information

Feature change description

From this release, you can view routes whose state is **Backup** when displaying detailed routing table information.

Command changes

Modified command: display ip routing-table

Syntax

```
display ip routing-table [ topology topo-name ] [ verbose ]  
display ip routing-table [ all-vpn-instance | vpn-instance  
vpn-instance-name ] [ verbose ]  
display ip routing-table [ all-routes ]
```

Views

Any view

Change description

Before modification: The command does not support displaying routes whose state is **Backup** in the detailed routing table information.

After modification: The command supports displaying routes whose state is **Backup** in the detailed routing table information. A backup route indicates a backup next hop calculated for ECMP routes.

Modified feature: Configuring the maximum number of Add-Path optimal routes that can be advertised to all peers

Feature change description

From this release, this feature is supported in BGP EVPN address family view.

Command changes

Modified command: additional-paths select-best

Syntax

```
additional-paths select-best best-number  
undo additional-paths select-best
```

Old views

- BGP IPv4 unicast address family view
- BGP-VPN IPv4 unicast address family view
- BGP VPNv4 address family view
- BGP IPv6 unicast address family view
- BGP-VPN IPv6 unicast address family view
- BGP VPNv6 address family view

New views

- BGP IPv4 unicast address family view
- BGP-VPN IPv4 unicast address family view
- BGP VPNv4 address family view
- BGP IPv6 unicast address family view
- BGP-VPN IPv6 unicast address family view
- BGP VPNv6 address family view
- BGP EVPN address family view

Change description

Before modification: This command is not supported in BGP EVPN address family view.

After modification: This command is supported in BGP EVPN address family view.

Modified feature: Configuring the BGP Additional Paths capabilities

Feature change description

From this release, this feature is supported in BGP EVPN address family view.

Command changes

Modified command: peer additional-paths

Syntax

In BGP VPNv4 address family view/BGP-VPN VPNv4 address family view/BGP VPNv6 address family view:

```
peer { group-name | ipv4-address [ mask-length ] } additional-paths  
{ receive | send } *
```

```
undo peer { group-name | ipv4-address [ mask-length ] } additional-paths  
{ receive | send } *
```

In BGP IPv4 unicast address family view/BGP-VPN IPv4 unicast address family view/BGP EVPN address family view:

```
peer { group-name | ipv4-address [ mask-length ] | ipv6-address  
[ prefix-length ] } additional-paths { receive | send } *
```

```
undo peer { group-name | ipv4-address [ mask-length ] | ipv6-address  
[ prefix-length ] } additional-paths { receive | send } *
```

In BGP IPv6 unicast address family view/BGP-VPN IPv6 unicast address family view:

```
peer { group-name | ipv4-address [ mask-length ] | ipv6-address  
[ prefix-length ] | link-local-address interface interface-type  
interface-number } additional-paths { receive | send } *
```

```
undo peer { group-name | ipv4-address [ mask-length ] | ipv6-address  
[ prefix-length ] | link-local-address interface interface-type  
interface-number } additional-paths { receive | send } *
```

Old views

BGP IPv4 unicast address family view

BGP-VPN IPv4 unicast address family view

BGP VPNv4 address family view

BGP-VPN VPNv4 address family view

BGP IPv6 unicast address family view

BGP-VPN IPv6 unicast address family view

BGP VPNv6 address family view

New views

BGP IPv4 unicast address family view

BGP-VPN IPv4 unicast address family view

BGP VPNv4 address family view

BGP IPv6 unicast address family view
BGP-VPN IPv6 unicast address family view
BGP VPNv6 address family view
BGP EVPN address family view

Change description

Before modification: This command is not supported in BGP EVPN address family view.

After modification: This command is supported in BGP EVPN address family view.

Modified feature: Configuring the maximum number of Add-Path optimal routes that can be advertised to a peer or peer group

Feature change description

From this release, this feature is supported in BGP EVPN address family view.

Command changes

Modified command: peer advertise additional-paths best

Syntax

In BGP VPNv4 address family view/BGP-VPN VPNv4 address family view/BGP VPNv6 address family view:

```
peer { group-name | ipv4-address [ mask-length ] } advertise additional-paths best number
```

```
undo peer { group-name | ipv4-address [ mask-length ] } advertise additional-paths best
```

In BGP IPv4 unicast address family view/BGP-VPN IPv4 unicast address family view/BGP EVPN address family view:

```
peer { group-name | ipv4-address [ mask-length ] | ipv6-address [ prefix-length ] } advertise additional-paths best number
```

```
undo peer { group-name | ipv4-address [ mask-length ] | ipv6-address [ prefix-length ] } advertise additional-paths best
```

In BGP IPv6 unicast address family view/BGP-VPN IPv6 unicast address family view:

```
peer { group-name | ipv4-address [ mask-length ] | ipv6-address [ prefix-length ] | link-local-address interface interface-type interface-number } advertise additional-paths best number
```

```
undo peer { group-name | ipv4-address [ mask-length ] | ipv6-address [ prefix-length ] | link-local-address interface interface-type interface-number } advertise additional-paths best
```

Old views

BGP IPv4 unicast address family view
BGP-VPN IPv4 unicast address family view

BGP VPNv4 address family view
BGP-VPN VPNv4 address family view
BGP IPv6 unicast address family view
BGP-VPN IPv6 unicast address family view
BGP VPNv6 address family view

New views

BGP IPv4 unicast address family view
BGP-VPN IPv4 unicast address family view
BGP VPNv4 address family view
BGP-VPN VPNv4 address family view
BGP IPv6 unicast address family view
BGP-VPN IPv6 unicast address family view
BGP VPNv6 address family view
BGP EVPN address family view

Change description

Before modification: This command is not supported in BGP EVPN address family view.

After modification: This command is supported in BGP EVPN address family view.

Modified feature: Password control restriction features

Feature change description

From this release, in non-FIPS, the following password restriction features are independent of the enabling status of the global password control feature:

- Minimum password length.
- Password composition policy.
- Username checking.

By default, in non-FIPS mode, the following settings apply:

- The minimum password length is 10 characters.
- A password must contain a minimum of two character types and a minimum of one character for each type.
- A password cannot contain the username or the reverse of the username.

If you downgrade the device from Release F2707 to an earlier version, the above password restriction configuration in Release F2707 will be saved in the configuration file as follows:

- If you have configured the above password restriction settings and have saved them by using the **save** command, the configured settings exist in the configuration file after the downgrading.
- If you have configured the above password restriction settings but have not saved them, the default password restriction settings in F2707 are in the configuration file after the downgrading.
- If you did not edit any password restriction settings in F2707, the default settings in F2707 are in the configuration file after the downgrading.

To restore the default settings of these three password restriction features in an earlier version, use the following commands:

- `password-control composition type-number 1 type-length 1`
- `undo password-control complexity user-name check`
- `password-control super composition type-number 1 type-length 1`

Command changes

Modified command: password-control composition

Syntax

```
password-control composition type-number type-number [ type-length
type-length ]
undo password-control composition
```

Views

System view
User group view
Local user view

Change description

Before modification: By default, in non-FIPS mode, a password must contain a minimum of one character type and a minimum of one character for each type.

After modification: By default, in non-FIPS mode, a password must contain a minimum of two character types and a minimum of one character for each type.

Modified command: password-control super composition

Syntax

```
password-control super composition type-number type-number [ type-length
type-length ]
undo password-control super composition
```

Views

System view

Change description

Before modification: By default, in non-FIPS mode, a super password must contain a minimum of one character type and a minimum of one character for each type.

After modification: By default, in non-FIPS mode, a super password must contain a minimum of two character types and a minimum of one character for each type.

Modified command: password-control complexity

Syntax

```
password-control complexity { same-character | user-name } check
undo password-control complexity
```

Views

System view
User group view

Local user view

Change description

Before modification: By default, in non-FIPS mode, username checking is disabled, and a password can contain the username and the reverse of the username.

After modification: By default, in non-FIPS mode, username checking is disabled, and a password cannot contain the username or the reverse of the username.

Modified feature: Flow-mirroring traffic to interfaces

Feature change description

From this release, you can configure flow-mirroring traffic to interfaces without specifying interfaces. Instead, you can specify encapsulation parameters for mirrored packets. When receiving packets that meet the flow mirroring criteria, the device performs the following tasks:

1. Encapsulates the packets based on the configured encapsulation parameters, including the source and destination IP addresses.
2. Searches the routing table for routes matching the configured destination IP address.
If ECMP load-balancing is configured, multiple matching ECMP routes might be found for the specified destination IP address.
3. Sends the mirrored packets out of the output interfaces in the matching routes to the final destination.

If a selected output interface goes down, the device will send the mirrored packets out of another interface based on the routing metric recalculations.

Command changes

Modified command: mirror-to interface

Old syntax

```
mirror-to interface interface-type interface-number [ loopback |  
[ destination-ip destination-ip-address source-ip source-ip-address  
[ dscp dscp-value | vlan vlan-id | vrf-instance vrf-name ] * ] ]  
undo mirror-to interface interface-type interface-number
```

New syntax

Syntax 1:

```
mirror-to interface interface-type interface-number [ loopback |  
[ destination-ip destination-ip-address source-ip source-ip-address  
[ dscp dscp-value | vlan vlan-id | vrf-instance vrf-name ] * ] ]  
undo mirror-to interface interface-type interface-number
```

Syntax 2:

```
mirror-to interface destination-ip destination-ip-address source-ip  
source-ip-address [ dscp dscp-value | vlan vlan-id | vrf-instance vrf-name ]  
*  
undo mirror-to interface destination-ip destination-ip-address source-ip  
source-ip-address
```

Views

Traffic behavior view

Change description

Before modification: Only syntax 1 is supported. You must specify a destination interface when configuring flow-mirroring traffic to an interface.

After modification: Both syntax 1 and syntax 2 are supported. When using syntax 2, you can configure flow-mirroring traffic to interfaces without specifying interfaces. The output interfaces of matching routes are the destination interfaces. When multiple routes are available for the specified destination IP address, mirrored packets can be flow-mirrored to multiple interfaces at the same time.

Modified feature: Configuring ARP scanning

Feature change description

From this version, you can set the ARP packet sending rate, in packets per minute (ppm).

Command changes

Modified command: arp scan

Old syntax

```
arp scan [ start-ip-address to end-ip-address ] [ send-rate pps ]
```

New syntax

```
arp scan [ start-ip-address to end-ip-address ] [ send-rate { ppm ppm |  
pps } ]
```

Views

Layer 3 Ethernet interface view

Layer 3 Ethernet subinterface view

Layer 3 aggregate interface view

Layer 3 aggregate subinterface view

VSI interface view

VLAN interface view

L3VE interface view

Change description

The **ppm ppm** option was added.

ppm ppm: Specifies the ARP packet sending rate, in packets per minute (ppm). The value range for the *ppm* argument is 10 to 600, and the value must be a multiple of 10.

Modified feature: NETCONF logging

Feature change description

This feature introduced the following changes:

- Added support for generating one log message for each row operation in an <action> or <edit-config> operation request.
- Changed the default setting for NETCONF logging.
 - **Before modification**—By default, NETCONF logging is disabled for all operations.
 - **After modification**—By default, NETCONF logging is enabled to log each row operation in <action> and <edit-config> operation requests and is disabled for all other types of operation requests.

Command changes

Modified command: netconf log

Old syntax

```
netconf log source { all | { agent | soap } * } { protocol-operation { all  
| { action | config | get | session | set | syntax | others } * } | verbose }  
undo netconf log source { all | { agent | soap } * } { protocol-operation  
{ all | { action | config | get | session | set | syntax | others } * }  
| verbose }
```

New syntax

```
netconf log source { all | { agent | soap } * } { protocol-operation { all  
| { action | config | get | session | set | syntax | others } * } | row-operation  
| verbose }  
undo netconf log source { all | { agent | soap } * } { protocol-operation  
{ all | { action | config | get | session | set | syntax | others } * }  
| row-operation | verbose }
```

Views

System view

Change description

- The **row-operation** keyword was added. This keyword enables the device to generate one log message for each row operation in an <action> and <edit-config> operation request.
- The default setting for NETCONF logging was changed.
 - **Before modification**—By default, NETCONF logging is disabled for all operations.
 - **After modification**—By default, NETCONF logging is enabled to log each row operation in <action> and <edit-config> operation requests and is disabled for all other types of operation requests.

Modified feature: Enabling route router MAC replacement for a peer or peer group on an EVPN-DCI ED

Feature change description

As from this version, you can specify the **dci** keyword when using the **peer router-mac-local** command.

Command changes

Modified command: **peer router-mac-local**

Old syntax

```
peer { group-name | ipv4-address [ mask-length ] | ipv6-address [ prefix-length ] } router-mac-local  
undo peer { group-name | ipv4-address [ mask-length ] | ipv6-address [ prefix-length ] } router-mac-local
```

New syntax

```
peer { group-name | ipv4-address [ mask-length ] | ipv6-address [ prefix-length ] } router-mac-local [ dci ]  
undo peer { group-name | ipv4-address [ mask-length ] | ipv6-address [ prefix-length ] } router-mac-local
```

Views

BGP EVPN address family view

Change description

Before modification: The **dci** keyword is not supported.

After modification: The **dci** keyword is supported.

dci: Establishes VXLAN-DCI tunnels with the peer or peer group. If you do not specify this keyword, whether the device establishes VXLAN-DCI tunnels with the peer or peer group depends on the **dci enable** command configuration in interface view.

Modified feature: Value range change for the extended community attribute list number in an RR reflection policy

Feature change description

In the **rr-filter** command, the value range for the *ext-comm-list-number* argument was changed to 1 to 65534.

Command changes

Modified command: **rr-filter**

Syntax

```
rr-filter ext-comm-list-number  
undo rr-filter
```

Change description

Before modification: The value range for the *ext-comm-list-number* argument is 1 to 199.

After modification: The value range for the *ext-comm-list-number* argument is 1 to 65534.

Modified feature: Allowing a DR interface to be up when it is the only member of a DR group

Feature change description

As of this software version, you can assign only one DR interface to a DR group and allow it to stay in up state.

Command changes

Modified command: **port drni group**

Old syntax

```
port drni group group-id  
undo port drni group
```

New syntax

```
port drni group group-id [ allow-single-member ]  
undo port drni group
```

Views

Layer 2 aggregate interface view

Parameters

allow-single-member: Allows the Layer 2 aggregate interface to be the only member of the specified DR group. DRNI MAD will not place the interface in DRNI MAD DOWN state. If you do not specify this keyword, DRNI MAD will shut down the interface if it is the only member of the specified DR group.

Change description

Before modification: DRNI MAD will shut down a Layer 2 aggregate interface if it is the only member of a DR group.

After modification: You can assign only one DR interface to a DR group and allow it to stay in up state.

Modified feature: Configuring a static route

Feature change description

As from this version, you can enable support for next hop recursion loops for an IPv4 static route.

Command changes

Modified command: ip route-static

Old syntax

```
ip route-static dest-address { mask-length | mask } interface-type
interface-number [ next-hop-address [ nexthop-index index-string ] ]
[ backup-interface interface-type interface-number [ backup-nexthop
backup-nexthop-address ] [ permanent ] | bfd { control-packet | echo-packet }
| permanent | track track-entry-number ] [ preference preference ] [ tag
tag-value ] [ description text ]
```

```
ip route-static dest-address { mask-length | mask } next-hop-address
[ nexthop-index index-string ] [ recursive-lookup host-route ] [ bfd
control-packet bfd-source ip-address | permanent | track
track-entry-number ] [ preference preference ] [ tag tag-value ]
[ description text ]
```

```
ip route-static dest-address { mask-length | mask } vpn-instance
d-vpn-instance-name next-hop-address [ nexthop-index index-string ]
[ recursive-lookup host-route ] [ bfd control-packet bfd-source ip-address
| permanent | track track-entry-number ] [ preference preference ] [ tag
tag-value ] [ description text ]
```

```
ip route-static group group-name interface-type interface-number
[ next-hop-address ] [ backup-interface interface-type interface-number
[ backup-nexthop backup-nexthop-address ] [ permanent ] | bfd
{ control-packet | echo-packet } | permanent | track track-entry-number ]
[ preference preference ] [ tag tag-value ] [ description text ]
```

```
ip route-static group group-name next-hop-address [ recursive-lookup
host-route ] [ bfd control-packet bfd-source ip-address | permanent | track
track-entry-number ] [ preference preference ] [ tag tag-value ]
[ description text ]
```

```
ip route-static group group-name vpn-instance d-vpn-instance-name
next-hop-address [ recursive-lookup host-route ] [ bfd control-packet
bfd-source ip-address | permanent | track track-entry-number ]
[ preference preference ] [ tag tag-value ] [ description text ]
```

```
undo ip route-static { dest-address { mask-length | mask } | group
group-name } [ interface-type interface-number [ next-hop-address ] |
next-hop-address | vpn-instance d-vpn-instance-name next-hop-address ]
[ preference preference ]
```

New syntax

```
ip route-static dest-address { mask-length | mask } interface-type
interface-number [ next-hop-address [ nexthop-index index-string ] ]
[ backup-interface interface-type interface-number [ backup-nexthop
backup-nexthop-address ] [ permanent ] | bfd { control-packet | echo-packet }
```

```

| permanent | track track-entry-number ] [ preference preference ] [ tag
tag-value ] [ recursive-lookup ] [ description text ]

ip route-static dest-address { mask-length | mask } next-hop-address
[ nexthop-index index-string ] [ recursive-lookup [ host-route ] ] [ bfd
control-packet bfd-source ip-address | permanent | track
track-entry-number ] [ preference preference ] [ tag tag-value ]
[ description text ]

ip route-static dest-address { mask-length | mask } vpn-instance
d-vpn-instance-name next-hop-address [ nexthop-index index-string ]
[ recursive-lookup [ host-route ] ] [ bfd control-packet bfd-source
ip-address | permanent | track track-entry-number ] [ preference
preference ] [ tag tag-value ] [ description text ]

ip route-static group group-name interface-type interface-number
[ next-hop-address ] [ backup-interface interface-type interface-number
[ backup-nexthop backup-nexthop-address ] [ permanent ] | bfd
{ control-packet | echo-packet } | permanent | track track-entry-number ]
[ preference preference ] [ tag tag-value ] [ recursive-lookup ]
[ description text ]

ip route-static group group-name next-hop-address [ recursive-lookup
[ host-route ] ] [ bfd control-packet bfd-source ip-address | permanent |
track track-entry-number ] [ preference preference ] [ tag tag-value ]
[ description text ]

ip route-static group group-name vpn-instance d-vpn-instance-name
next-hop-address [ recursive-lookup [ host-route ] ] [ bfd control-packet
bfd-source ip-address | permanent | track track-entry-number ]
[ preference preference ] [ tag tag-value ] [ description text ]

undo ip route-static { dest-address { mask-length | mask } | group
group-name } [ interface-type interface-number [ next-hop-address ] |
next-hop-address | vpn-instance d-vpn-instance-name next-hop-address ]
[ preference preference ]

```

Views

System view

Change description

Before modification: The **recursive-lookup** keyword was not supported.

After modification: The **recursive-lookup** keyword is supported. This keyword enables support for next hop recursion loops for the IPv4 static route. If you do not specify this keyword, the device does not support next hop recursion loops for the IPv4 static route.

Modified command: ip route-static vpn-instance

Old syntax

```

ip route-static vpn-instance s-vpn-instance-name dest-address
{ mask-length | mask } interface-type interface-number [ next-hop-address
[ nexthop-index index-string ] ] [ backup-interface interface-type
interface-number [ backup-nexthop backup-nexthop-address ] [ permanent ] |
bfd { control-packet | echo-packet } | permanent | track track-entry-number ]
[ preference preference ] [ tag tag-value ] [ description text ]

ip route-static vpn-instance s-vpn-instance-name dest-address
{ mask-length | mask } vpn-instance d-vpn-instance-name next-hop-address
[ nexthop-index index-string ] [ recursive-lookup host-route ] [ bfd

```

```

control-packet  bfd-source  ip-address  |  permanent  |  track
track-entry-number  ] [  preference  preference  ] [  tag  tag-value  ]
[ description text ]

ip  route-static  vpn-instance  s-vpn-instance-name  dest-address
{ mask-length | mask } next-hop-address [ nexthop-index index-string ]
[ public ] [ bfd control-packet bfd-source ip-address | permanent | track
track-entry-number  ] [  preference  preference  ] [  tag  tag-value  ]
[ description text ]

ip  route-static  vpn-instance  s-vpn-instance-name  dest-address
{ mask-length | mask } vpn-instance d-vpn-instance-name next-hop-address
[ nexthop-index index-string ] [ recursive-lookup host-route ] [ bfd
control-packet  bfd-source  ip-address  |  permanent  |  track
track-entry-number  ] [  preference  preference  ] [  tag  tag-value  ]
[ description text ]

ip  route-static  vpn-instance  s-vpn-instance-name  dest-address
{ mask-length | mask } next-hop-address [ nexthop-index index-string ]
[ recursive-lookup host-route ] [ public ] [ bfd control-packet bfd-source
ip-address ] [ preference preference ] [ tag tag-value ] [ description text ]

ip  route-static  vpn-instance  s-vpn-instance-name  group  group-name
interface-type  interface-number  [  next-hop-address  ] [  bfd
{ control-packet | echo-packet } | backup-interface interface-type
interface-number [ backup-nexthop backup-nexthop-address ] [ permanent ] ]
[ preference preference ] [ tag tag-value ] [ description text ]

ip  route-static  vpn-instance  s-vpn-instance-name  group  group-name
next-hop-address [ recursive-lookup host-route ] [ public ] [ bfd
control-packet  bfd-source  ip-address  |  permanent  |  track
track-entry-number  ] [  preference  preference  ] [  tag  tag-value  ]
[ description text ]

undo ip route-static vpn-instance s-vpn-instance-name { dest-address
{ mask-length | mask } | group group-name } [ interface-type
interface-number [ next-hop-address ] | next-hop-address [ public ] |
vpn-instance d-vpn-instance-name next-hop-address ] [ preference
preference ]

```

New syntax

```

ip  route-static  vpn-instance  s-vpn-instance-name  dest-address
{ mask-length | mask } interface-type interface-number [ next-hop-address
[ nexthop-index index-string ] ] [ backup-interface interface-type
interface-number [ backup-nexthop backup-nexthop-address ] [ permanent ] |
bfd { control-packet | echo-packet } | permanent | track track-entry-number ]
[ preference preference ] [ tag tag-value ] [ recursive-lookup ]
[ description text ]

ip  route-static  vpn-instance  s-vpn-instance-name  dest-address
{ mask-length | mask } vpn-instance d-vpn-instance-name next-hop-address
[ nexthop-index index-string ] [ recursive-lookup [ host-route ] ] [ bfd
control-packet  bfd-source  ip-address  |  permanent  |  track
track-entry-number  ] [  preference  preference  ] [  tag  tag-value  ]
[ description text ]

ip  route-static  vpn-instance  s-vpn-instance-name  dest-address
{ mask-length | mask } vpn-instance d-vpn-instance-name next-hop-address
[ nexthop-index index-string ] [ recursive-lookup [ host-route ] ] [ bfd
control-packet  bfd-source  ip-address  |  permanent  |  track

```

```

track-entry-number ] [ preference preference ] [ tag tag-value ]
[ description text ]

ip route-static vpn-instance s-vpn-instance-name dest-address
{ mask-length | mask } next-hop-address [ nexthop-index index-string ]
[ recursive-lookup [ host-route ] ] [ public ] [ bfd control-packet
bfd-source ip-address | permanent | track track-entry-number ]
[ preference preference ] [ tag tag-value ] [ description text ]

ip route-static vpn-instance s-vpn-instance-name group group-name
interface-type interface-number [ next-hop-address ] [ bfd
{ control-packet | echo-packet } | backup-interface interface-type
interface-number [ backup-nexthop backup-nexthop-address ] [ permanent ] ]
[ preference preference ] [ tag tag-value ] [ recursive-lookup ]
[ description text ]

ip route-static vpn-instance s-vpn-instance-name group group-name
next-hop-address [ recursive-lookup [ host-route ] ] [ public ] [ bfd
control-packet bfd-source ip-address | permanent | track
track-entry-number ] [ preference preference ] [ tag tag-value ]
[ description text ]

undo ip route-static vpn-instance s-vpn-instance-name { dest-address
{ mask-length | mask } | group group-name } [ interface-type
interface-number [ next-hop-address ] | next-hop-address [ public ] |
vpn-instance d-vpn-instance-name next-hop-address ] [ preference
preference ]

```

Views

System view

Change description

Before modification: The **recursive-lookup** keyword was not supported.

After modification: The **recursive-lookup** keyword is supported. This keyword enables support for next hop recursion loops for the IPv4 static route. If you do not specify this keyword, the device does not support next hop recursion loops for the IPv4 static route.

Modified feature: Configuring an IPv6 static route

Feature change description

As from this version, you can enable support for next hop recursion loops for an IPv6 static route.

Command changes

Modified command: ipv6 route-static

Old syntax

```

ipv6 route-static ipv6-address prefix-length { interface-type
interface-number [ next-hop-address ] [ bfd { control-packet | echo-packet }
[ bfd-source ipv6-address ] | permanent | track track-entry-number ] |
[ vpn-instance d-vpn-instance-name ] next-hop-address [ bfd
control-packet bfd-source ipv6-address | permanent | track
track-entry-number ] } [ preference preference ] [ tag tag-value ]
[ description text ]

```

```

undo ipv6 route-static ipv6-address prefix-length [ interface-type
interface-number [ next-hop-address ] | [ vpn-instance
d-vpn-instance-name ] next-hop-address ] [ preference preference ]

ipv6 route-static vpn-instance s-vpn-instance-name ipv6-address
prefix-length { interface-type interface-number [ next-hop-address ] [ bfd
{ control-packet | echo-packet } [ bfd-source ipv6-address ] | permanent |
track track-entry-number ] | next-hop-address [ public ] [ bfd
control-packet bfd-source ipv6-address | permanent | track
track-entry-number ] | vpn-instance d-vpn-instance-name next-hop-address
[ bfd control-packet bfd-source ipv6-address | permanent | track
track-entry-number ] } [ preference preference ] [ tag tag-value ]
[ description text ]

undo ipv6 route-static vpn-instance s-vpn-instance-name ipv6-address
prefix-length [ interface-type interface-number [ next-hop-address ] |
next-hop-address [ public ] | vpn-instance d-vpn-instance-name
next-hop-address ] [ preference preference ]

```

New syntax

```

ipv6 route-static ipv6-address prefix-length interface-type
interface-number [ next-hop-address ] [ bfd { control-packet | echo-packet }
[ bfd-source ipv6-address ] | permanent | track track-entry-number ]
[ preference preference ] [ tag tag-value ] [ description text |
recursive-lookup ]

ipv6 route-static ipv6-address prefix-length [ vpn-instance
d-vpn-instance-name ] next-hop-address [ bfd control-packet bfd-source
ipv6-address | permanent | track track-entry-number ] [ preference
preference ] [ tag tag-value ] [ description text | recursive-lookup ]

undo ipv6 route-static ipv6-address prefix-length [ interface-type
interface-number [ next-hop-address ] | [ vpn-instance
d-vpn-instance-name ] next-hop-address ] [ preference preference ]

ipv6 route-static vpn-instance s-vpn-instance-name ipv6-address
prefix-length interface-type interface-number [ next-hop-address ] [ bfd
{ control-packet | echo-packet [ bfd-source ipv6-address ] | permanent |
track track-entry-number ] } [ preference preference ] [ tag tag-value ]
[ description text | recursive-lookup ]

ipv6 route-static vpn-instance s-vpn-instance-name ipv6-address
prefix-length { next-hop-address [ public ] [ bfd control-packet bfd-source
ipv6-address | permanent | track track-entry-number ] | vpn-instance
d-vpn-instance-name next-hop-address [ bfd control-packet bfd-source
ipv6-address | permanent | track track-entry-number ] } [ preference
preference ] [ tag tag-value ] [ description text | recursive-lookup ]

undo ipv6 route-static vpn-instance s-vpn-instance-name ipv6-address
prefix-length [ interface-type interface-number [ next-hop-address ] |
next-hop-address [ public ] | vpn-instance d-vpn-instance-name
next-hop-address ] [ preference preference ]

```

Views

System view

Change description

Before modification: The **recursive-lookup** keyword was not supported.

After modification: The **recursive-lookup** keyword is supported. This keyword enables support for next hop recursion loops for the IPv6 static route. If you do not specify this keyword, the device does not support next hop recursion loops for the IPv6 static route.

Modified feature: Displaying OSPFv3 LSDB information

Feature change description

As from this version, the **display ospfv3 lsdb** command can display Type-33, Type-35, and Type-36 LSAs.

Command changes

Modified command: **display ospfv3 lsdb**

Old syntax

```
display ospfv3 [ process-id ] lsdb [ { external | grace | inter-prefix |  
inter-router | intra-prefix | link | network | nssa | router | unknown  
[ type ] } [ link-state-id ] [ originate-router router-id | self-originate ]  
[ statistics | total | verbose ]
```

New syntax

```
display ospfv3 [ process-id ] lsdb [ { e-inter-prefix | e-inter-router |  
e-router | external | grace | inter-prefix | inter-router | intra-prefix  
| link | network | nssa | router | unknown [ type ] } [ link-state-id ]  
[ originate-router router-id | self-originate ] [ statistics | total |  
verbose ]
```

Change description

Before modification: This command cannot display Type-33, Type-35, or Type-36 LSAs.

After modification: The **e-inter-prefix**, **e-inter-router**, and **e-router** keywords were added to this command. You can specify these keywords to display Type-33, Type-35, and Type-36 LSAs.

Modified feature: Configuring BGP extended community lists

Feature change description

As from this version, the supported value range for the extended community list number changes to 1 to 65534.

Command changes

Modified command: **display ip extcommunity-list**

Use **display ip extcommunity-list** to display BGP extended community list information.

Syntax

```
display ip extcommunity-list [ ext-comm-list-number ]
```

Views

Any view

Change description

Before modification: The value range for the *ext-comm-list-number* argument was 1 to 199.

After modification: The value range for the *ext-comm-list-number* argument is 1 to 65534.

Modified command: if-match extcommunity

Use **if-match extcommunity** to match BGP routes whose extended community attribute matches a specified extended community list.

Use **undo if-match extcommunity** to remove the specified extended community list match criterion.

Syntax

```
if-match extcommunity ext-comm-list-number&<1-32>
```

```
undo if-match extcommunity [ ext-comm-list-number&<1-32> ]
```

Views

Routing policy node view

Change description

Before modification: The value range for the *ext-comm-list-number* argument was 1 to 199.

After modification: The value range for the *ext-comm-list-number* argument is 1 to 65534.

Modified command: ip extcommunity-list

Use **ip extcommunity-list** to configure an extended community list.

Use **undo ip extcommunity-list** to remove an extended community list.

Syntax

```
ip extcommunity-list ext-comm-list-number { deny | permit } { rt  
route-target | soo site-of-origin }&<1-32>
```

```
undo ip extcommunity-list ext-comm-list-number [ { deny | permit } [ rt  
route-target | soo site-of-origin ]&<1-32> ]
```

Views

System view

Change description

Before modification: The value range for the *ext-comm-list-number* argument was 1 to 199.

After modification: The value range for the *ext-comm-list-number* argument is 1 to 65534.

Modified feature: Replicating VLINK direct routes

Feature change description

As from this version, you can configure the public instance or a VPN instance to replicate VLINK direct routes.

The `route-replicate from vpn-instance protocol direct` or `route-replicate from public protocol direct` command does not replicate VLINK direct routes, which will cause traffic forwarding failures. To address this issue, you can specify the `vlink-direct` keyword to replicate VLINK direct routes.

Command changes

Modified command: route-replicate

Old syntax

In public instance IPv4 address family view:

```
route-replicate from vpn-instance vpn-instance-name protocol { direct |
static | bgp as-number | { isis | ospf | rip } process-id } [ advertise ]
[ route-policy route-policy-name ]
```

```
undo route-replicate from vpn-instance vpn-instance-name protocol
{ direct | static | bgp as-number | { isis | ospf | rip } process-id }
```

In VPN instance IPv4 address family view:

```
route-replicate from { public | vpn-instance vpn-instance-name } protocol
{ direct | static | bgp as-number | { isis | ospf | rip } process-id }
[ advertise ] [ route-policy route-policy-name ]
```

```
undo route-replicate from { public | vpn-instance vpn-instance-name }
protocol { direct | static | bgp as-number | { isis | ospf | rip } process-id }
```

In public instance IPv6 address family view:

```
route-replicate from vpn-instance vpn-instance-name protocol { direct |
static | bgp4+ as-number | { isisv6 | ospfv3 | ripng } process-id }
[ advertise ] [ route-policy route-policy-name ]
```

```
undo route-replicate from vpn-instance vpn-instance-name protocol
{ direct | static | bgp4+ as-number | { isisv6 | ospfv3 | ripng } process-id }
```

In VPN instance IPv6 address family view:

```
route-replicate from { public | vpn-instance vpn-instance-name } protocol
{ direct | static | bgp4+ as-number | { isisv6 | ospfv3 | ripng } process-id }
[ advertise ] [ route-policy route-policy-name ]
```

```
undo route-replicate from { public | vpn-instance vpn-instance-name }
protocol { direct | static | bgp4+ as-number | { isisv6 | ospfv3 | ripng }
process-id }
```

New syntax

In public instance IPv4 address family view:

```
route-replicate from vpn-instance vpn-instance-name protocol { bgp
as-number | direct | static | vlink-direct | { isis | ospf | rip } process-id }
[ advertise ] [ route-policy route-policy-name ]
```

```
undo route-replicate from vpn-instance vpn-instance-name protocol { bgp
as-number | direct | static | vlink-direct | { isis | ospf | rip } process-id }
```

In VPN instance IPv4 address family view:

```
route-replicate from { public | vpn-instance vpn-instance-name } protocol
{ bgp as-number | direct | static | vlink-direct | { isis | ospf | rip }
process-id } [ advertise ] [ route-policy route-policy-name ]
```

```
undo route-replicate from { public | vpn-instance vpn-instance-name }
protocol { bgp as-number | direct | static | vlink-direct | { isis | ospf | rip }
process-id }
```

In public instance IPv6 address family view:

```
route-replicate from vpn-instance vpn-instance-name protocol { bgp4+
as-number | direct | static | vlink-direct | { isisv6 | ospfv3 | ripng }
process-id } [ advertise ] [ route-policy route-policy-name ]
```

```
undo route-replicate from vpn-instance vpn-instance-name protocol { bgp4+
as-number | direct | static | vlink-direct | { isisv6 | ospfv3 | ripng }
process-id }
```

In VPN instance IPv6 address family view:

```
route-replicate from { public | vpn-instance vpn-instance-name } protocol
{ bgp4+ as-number | direct | static | vlink-direct | { isisv6 | ospfv3 | ripng }
process-id } [ advertise ] [ route-policy route-policy-name ]
```

```
undo route-replicate from { public | vpn-instance vpn-instance-name }
protocol { bgp4+ as-number | direct | static | vlink-direct | { isisv6 |
ospfv3 | ripng } process-id }
```

Views

Public instance IPv4 address family view

VPN instance IPv4 address family view

Public instance IPv6 address family view

VPN instance IPv6 address family view

Change description

Before modification: The public instance and VPN instances cannot replicate VLINK direct routes.

After modification: The **vlink-direct** keyword was added to the **route-replicate** command. You can specify the **vlink-direct** keyword to replicate VLINK direct routes.

Modified feature: Specifying a controller for an OpenFlow switch and configuring the main connection to the controller

Feature change description

As from this release, you can specify a certificate-based access control policy when you configure the main connection to a controller.

Command changes

Modified command: controller address

Old syntax

```
controller controller-id address { ip ipv4-address | ipv6 ipv6-address }  
[ port port-number ] [ local address { ip local-ipv4-address | ipv6  
local-ipv6-address } [ port local-port- number ] ] [ ssl ssl-policy-name ]  
[ vrf vrf-name ]
```

New syntax

```
controller controller-id address { ip ipv4-address | ipv6 ipv6-address }  
[ port port-number ] [ local address { ip local-ipv4-address | ipv6  
local-ipv6-address } [ port local-port- number ] ] [ ssl ssl-policy-name  
[ access-control-policy acp-policy-name ] ] [ vrf vrf-name ]
```

Views

OpenFlow instance view

Change description

The **access-control-policy acp-policy-name** option was added.

access-control-policy acp-policy-name Specifies a certificate-based access control policy to verify the controller (SSL server) certificate. The *ssl-policy-name* argument is a case-insensitive string of 1 to 31 characters. Certificate-based access control policies are configured by using the **pki certificate access-control-policy** command. For more information about certificate-based access control policies, see PKI configuration in *Security Configuration Guide*.

Modified feature: Displaying the CPU usage statistics

Feature change description

As from this release, the device displays two hyphens (--) when the **display cpu-usage** command fails to obtain data.

Command changes

Modified command: display cpu-usage

Syntax

```
display cpu-usage [ summary ] [ slot slot-number [ cpu cpu-number [ core  
{ core-number | all } ] ] ]
```

Views

Any view

Usage guidelines

If two hyphens (--) are displayed for CPU usage values during the most recent 5-second, 1-minute, or 5-minute interval, the command might fail to obtain data from the device. You can execute the command again later.

Change description

Before modification: No information is displayed when the **display cpu-usage** command fails to obtain data.

After modification: The device displays two hyphens (--) when the **display cpu-usage** command fails to obtain data.

Modified feature: Displaying memory usage information

Feature change description

As from this release, the device displays two hyphens (--) when the **display memory** command fails to obtain data.

Command changes

Modified command: display memory

Syntax

```
display memory [ summary ] [ slot slot-number [ cpu cpu-number ] ]
```

Views

Any view

Usage guidelines

If two hyphens (--) are displayed for memory usage values during the most recent 5-second, 1-minute, or 5-minute interval, the command might fail to obtain data from the device. You can execute the command again later.

Change description

Before modification: No information is displayed when the **display memory** command fails to obtain data.

After modification: The device displays two hyphens (--) when the **display memory** command fails to obtain data.

Modified feature: Password complexity checking

Feature change description

In earlier releases, if password authentication is enabled, and the login password is weak, the system forces the user to change it. The user can log in to the device only after the password meets the password requirements.

In this release, when password authentication is enabled, the following rules apply:

- If the default password is used for login, the system forces the user to change it. The user can log in to the device only after the password meets the password requirements.
- If a non-default password is used for login and the password is weak, the system displays a message about a weak password but does not force the user to change it. As a best practice to improve the device security, the user can change the identified weak password to meet the password requirements.

Command changes

None.

Modified feature: Configuring kernel thread deadlock detection

Feature change description

As from this release, upon detecting a thread that occupies the CPU for a specific interval, the device determines that a deadlock has occurred, logs the event, and reboots to remove the deadlock.

Command changes

Modified command: monitor kernel deadlock enable

Syntax

```
monitor kernel deadlock enable [ slot slot-number [ cpu cpu-number [ core core-number&<1-64> ] ] ]  
undo monitor kernel deadlock enable [ slot slot-number [ cpu cpu-number ] ]
```

Views

System view

Change description

Before modification: This feature enables the device to detect deadlocks. If a thread occupies the CPU for a specific interval, the device determines that a deadlock has occurred and logs the event.

After modification: This feature enables the device to detect deadlocks. If a thread occupies the CPU for a specific interval, the device determines that a deadlock has occurred, logs the event, and reboots to remove the deadlock.

Modified feature: Configuring a traffic redirecting action in a QoS traffic behavior

Feature change description

As from this release, the action of redirecting traffic to a tunnel group is supported.

Command changes

Modified command: redirect

Old syntax

```
redirect { cpu | interface interface-type interface-number }  
undo redirect { cpu | interface interface-type interface-number }
```

New syntax

```
redirect { cpu | interface interface-type interface-number | tunnel-group
group-id }

undo redirect { cpu | interface interface-type interface-number |
tunnel-group group-id }
```

Views

Traffic behavior view

Parameters

tunnel-group *group-id*: Specifies a tunnel group by its ID in the range of 1 to 32. For more information about tunnel groups.

Change description

Before modification: The action of redirecting traffic to a tunnel group is not supported.

After modification: The action of redirecting traffic to a tunnel group is supported.

Modified feature: DRNI IPP configuration

Feature change description

Before modification: The link type and VLAN assignment settings are not automatically issued to an aggregate interface after you configure it as the IPP. You can assign the IPP role to an aggregate interface before you assign member ports to the corresponding aggregation group.

After modification: The link type and VLAN assignment settings are automatically issued to an aggregate interface after you configure it as the IPP. Before you assign the IPP role to an aggregate interface, you must assign member ports to the corresponding aggregation group.

[Sysname-Bridge-Aggregation11] display this

```
#
interface Bridge-Aggregation11
port link-type trunk
port trunk permit vlan all
port drni intra-portal-port 1
#
return
```

Command changes

None.

Release 2702

This release has the following changes:

- New features: Fundamentals features
- New features: Layer 2—LAN switching features
- New features: Layer 3—IP services features
- New features: Layer 3—IP routing features
- New features: IP multicast features
- New features: MPLS features
- New features: ACL and QoS features
- New features: Security features
- New features: High availability features
- New features: OpenFlow features
- New features: VXLAN features
- New features: EVPN features
- Modified feature: Software patching by using `issu` commands
- Modified feature: Automatic configuration
- Modified feature: Collision handling in BFD MAD
- Modified feature: Processing after the link mode of an Ethernet interface is switched
- Modified feature: Link state change suppression on an interface
- Modified feature: Setting the global aging timer for ND entries in stale state
- Modified feature: Setting the interface-specific aging timer for ND entries in stale state
- Modified feature: Configuring OSPF area authentication
- Modified feature: Configuring OSPF interface authentication
- Modified feature: Configuring a virtual link
- Modified feature: Displaying statistics for the IGMP messages and PIMv2 hello messages learned through IGMP snooping
- Modified feature: Displaying statistics for the MLD messages and IPv6 PIM hello messages learned through MLD snooping
- Modified feature: Enabling a VPN instance to replicate routes from the public network or other VPN instances
- Modified feature: Using BFD to verify LSP or MPLS TE tunnel connectivity
- Modified feature: MAC authentication VLAN mode
- Modified feature: Port security MAC move
- Modified feature: Creating a BFD session for detecting the local interface state
- Modified feature: Associating Track with application modules
- Modified feature: Configuring an interface event for a CLI-defined EAA monitor policy
- Modified feature: Configuring an EAA monitor policy by using `Tcl`
- Modified feature: Removing a TCP or UDP listening service for a VPN instance
- Modified feature: Specifying the role of the device in the VCF fabric
- Modified feature: Frame match criteria of VXLAN Ethernet service instances

- Modified feature: NETCONF logging

New features: Fundamentals features

Table 1 describes the Fundamentals features added in this software version. For more information about the features and commands, see *HPE FlexFabric 5940 & 5930 Switch Series Fundamentals Configuration Guide-Release 27xx* and *HPE FlexFabric 5940 & 5930 Switch Series Fundamentals Command Reference-Release 27xx*.

Table 1 Fundamentals features added in version R2702

| Feature | Command changes |
|---|-----------------|
| Python: Comware extended Python API channel | N/A |
| Python: Comware extended Python API send | N/A |
| Python: Comware extended Python API SYSLOG | N/A |

New features: Layer 2—LAN switching features

Table 2 describes the Layer 2—LAN switching features added in this software version. For more information about the features and commands, see *HPE FlexFabric 5940 & 5930 Switch Series Layer 2—LAN Switching-Release 27xx* and *HPE FlexFabric 5940 & 5930 Switch Series Layer 2—LAN Switching Command Reference-Release 27xx*.

Table 2 Layer 2—LAN switching features added in version R2702

| Feature | Command changes |
|--|---|
| Ethernet interfaces: Configuring PFC deadlock detection | <p>The following commands were added:</p> <ul style="list-style-type: none"> • <code>priority-flow-control deadlock cos</code> • <code>priority-flow-control deadlock auto-recover action</code> • <code>priority-flow-control deadlock auto-recover cos</code> • <code>priority-flow-control deadlock precision</code> • <code>priority-flow-control deadlock recover</code> • <code>priority-flow-control deadlock recover-mode</code> • <code>priority-flow-control deadlock enable</code> • <code>priority-flow-control deadlock threshold</code> |
| Ethernet link aggregation: Setting the global resilient link-aggregation load sharing mode. | The resilient keyword was added to the <code>link-aggregation global load-sharing mode</code> command. |
| Ethernet link aggregation: Configuring load sharing for tunneled traffic based on the inner IP header and outer IP header. | The all keyword was added to the <code>link-aggregation global load-sharing tunnel</code> command. |

| | |
|---|--|
| Ethernet link aggregation: Setting the resilient link-aggregation load sharing mode for an aggregation group. | The resilient keyword was added to the link-aggregation load-sharing mode command. |
| Ethernet link aggregation: Setting the minimum percentage of Selected ports in an aggregation group | The percentage <i>number</i> option was added to the link-aggregation selected-port minimum command. |
| Ethernet link aggregation: Setting the physical state change suppression interval on an aggregate interface | The link-delay command was added. |
| VLAN mapping: Configuring two-to-one VLAN mapping | vlan mapping egress |
| LLDP: Setting the timeout for receiving LLDP frames and enable the device to report no LLDP neighbor events | The lldp timer rx-timeout command was added. |
| LLDP: Clearing LLDP statistics on interfaces | The reset lldp statistics command was added. |

New features: Layer 3—IP services features

Table 3 describes the Layer 3—IP services features added in this software version. For more information about the features and commands, see *HPE FlexFabric 5940 & 5930 Switch Series Layer 3—IP Services-Release 27xx* and *HPE FlexFabric 5940 & 5930 Switch Series Layer 3—IP Services Command Reference-Release 27xx*.

Table 3 Layer 3—IP services features added in version R2702

| Feature | Command changes |
|---|--|
| DHCP: Enabling the DHCP server to return a DHCP-NAK message if the client notions of their IP addresses are incorrect | The dhcp server request-ip-address check command was added. |
| DHCP: Enabling the DHCP relay agent to insert Option 60 into DHCP requests | The dhcp relay insert option60 command was added. |
| IP forwarding basics: Enabling split horizon forwarding. | The forwarding split-horizon command was added. |
| IP performance optimization: Enabling the device to encapsulate the TCP Timestamps option in outgoing TCP packets | The tcp timestamps enable command was added. |
| IPv6 basics: Enabling ND direct route advertisement | The ipv6 nd route-direct advertise command was added. |

New features: Layer 3—IP routing features

Table 4 describes the Layer 3—IP routing features added in this software version. For more information about the features and commands, see *HPE FlexFabric 5940 & 5930 Switch Series Layer 3—IP Routing Configuration Guide-Release 27xx* and *HPE FlexFabric 5940 & 5930 Switch Series Layer 3—IP Routing Command Reference-Release 27xx*.

Table 4 Layer 3—IP routing features added in version R2702

| Feature | Command changes |
|--|--|
| OSPF: Displaying OSPF neighbor state change log information for the specified neighbor on the specified IRF member device. | The <i>neighbor-id</i> argument and slot slot-number option were added to the display ospf event-log command. |
| OSPF: Displaying information about the hello packets sent to or received from neighbors. | The hello keyword was added to the display ospf peer command. |
| OSPF: Displaying statistics for sent or received hello packets. | The hello keyword was added to the display ospf statistics command. |
| OSPF: Displaying log information about received or sent hello packets. | The display ospf event-log hello command was added. |
| OSPF: Displaying information about hello packets sent by OSPF interfaces. | The display ospf interface hello command was added. |
| OSPF: Setting the number of logs for received hello packets, abnormal hello packets received, received hello packets that were dropped, sent hello packets, abnormal hello packets sent, and hello packets that failed to be sent. | The hello , received , sent , abnormal , dropped , and failed keywords were added to the event-log command. |
| OSPF: Clearing OSPF neighbor state change information about an IRF member device. | The slot keyword was added to the reset ospf event-log command. |
| OSPF: Clearing OSPF log information about received or sent hello packets. | The reset ospf event-log hello command was added. |
| BGP: Enabling BGP to prefer routes with an IPv6 next hop during optimal route selection. | The bestroute ipv6-next-hop command was added. |
| BGP: Displaying information about dynamic IPv6 peers or peer groups for the EVPN address family and resetting or soft resetting the BGP session to an IPv6 peer for the EVPN address family. | The <i>ipv6-address</i> argument was added to the following commands: <ul style="list-style-type: none"> display bgp peer l2vpn evpn refresh bgp l2vpn evpn reset bgp l2vpn evpn |
| BGP: Displaying route flap statistics for BGP routes that match an AS path list specified by its name. | The <i>as-path-acl-name</i> argument was added to the display bgp routing-table flap-info command. |
| BGP: Displaying BGP IPv4 unicast routes, IPv4 multicast routes, IPv6 unicast routes, and IPv6 multicast routes that match an AS path list specified by its name. | The <i>as-path-acl-name</i> argument was added to the following commands: <ul style="list-style-type: none"> display bgp routing-table ipv4 unicast display bgp routing-table ipv4 multicast display bgp routing-table ipv6 unicast display bgp routing-table ipv6 multicast |
| BGP: Configuring BGP features for IPv6 peers in EVPN address family view. | The <i>ipv6-address</i> argument was added to the following commands in BGP EVPN address family view: <ul style="list-style-type: none"> peer enable peer next-hop-local |

| Feature | Command changes |
|--|--|
| | <ul style="list-style-type: none"> • peer reflect-client • peer route-policy • peer advertise-community • peer allow-as-loop |
| BGP: Specifying an AS path list by its name to filter routes incoming from or outgoing to a peer or peer group. | The <i>as-path-acl-name</i> argument was added to the peer as-path-acl command. |
| BGP: Displaying BGP update group information for IPv6 EVPN address family and displaying log information for an IPv6 EVPN peer. | <p>The <i>ipv6-address</i> argument was added to the following commands:</p> <ul style="list-style-type: none"> • display bgp update-group l2vpn evpn • display bgp peer l2vpn evpn log-info |
| BGP: Enabling the route reflector to change the attributes of routes to be reflected. | The reflect change-path-attribute command was added. |
| BGP: Specifying existent and nonexistent policies to control route advertisement. | <p>The following commands were added:</p> <ul style="list-style-type: none"> • peer advertise-policy exist-policy • peer advertise-policy non-exist-policy |
| BGP: Enabling BGP to ignore the first AS number of EBGP route updates received from a peer or peer group. | The peer ignore-first-as command was added. |
| BGP: Clearing flap statistics for BGP routes that match an AS path list specified by its name. | The <i>as-path-acl-name</i> argument was added to the reset bgp flap-info command. |
| Policy-based routing: Setting a local QoS ID match criterion for IP packets. | The if-match qos-local-id local-id-value [qppb-manipulation] command was added. |
| Policy-based routing: Specifying NULL0 as the output interface for IP packets. | The apply output-interface NULL0 command was added. |
| IPv6 policy-based routing: Setting a local QoS ID match criterion for IPv6 packets. | The if-match qos-local-id local-id-value [qppb-manipulation] command was added. |
| IPv6 policy-based routing: Specifying NULL0 as the output interface for IPv6 packets. | The apply output-interface NULL0 command was added. |
| Routing policy: Matching BGP routes whose AS_PATH attribute matches an AS path list specified by its name, configuring an AS path list and specifying its name, and displaying information about a BGP AS path list specified by its name. | <p>The <i>as-path-name</i> argument was added to the following commands:</p> <ul style="list-style-type: none"> • if-match as-path • ip as-path • display ip as-path |
| Routing policy: Clearing BGP AS path list statistics. | The reset ip as-path command was added. |
| Routing policy: Configuring an RD list, matching routes whose RD matches the specified RD list, and displaying RD list information. | <p>The following commands were added:</p> <ul style="list-style-type: none"> • ip rd-list • if-match rd-list • display ip rd-list |
| Routing policy: Matching IPv6 prefixes with the specified length whose last bit is an even number. | The if-match ipv6 even-prefix-length command was added. |

| Feature | Command changes |
|---|---|
| Routing policy: Matching IPv6 prefixes with the specified length whose last bit is an odd number. | The if-match ipv6 odd-prefix-length command was added. |

New features: IP multicast features

[Table 5](#) describes the IP multicast features added in this software version. For more information about the features and commands, see *HPE FlexFabric 5940 & 5930 Switch Series IP Multicast Configuration Guide-Release 27xx* and *HPE FlexFabric 5940 & 5930 Switch Series IP Multicast Command Reference-Release 27xx*.

Table 5 IP multicast features added in version R2702

| Feature | Command changes |
|---|--|
| IGMP snooping: Displaying multicast group information that IGMP snooping learns from EVPN | The display igmp-snooping evpn-group command was added. |
| PIM: Displaying PIM routing entries for MVPN extranet | The extranet { source-vpn-instance source-vpn-instance-name source-public-instance receive-vpn-instance receive-vpn-instance-name receive-public-instance } option was added to the display pim routing-table command. |
| Multicast VPN: Configuring an IPv4 MVPN extranet RPF selection policy | The multicast extranet select-rpf command was added. |
| Multicast VPN: Configuring an IPv6 MVPN extranet RPF selection policy | The ipv6 multicast extranet select-rpf command was added. |

New features: MPLS features

[Table 6](#) describes the MPLS features added in this software version. For more information about the features and commands, see *HPE FlexFabric 5940 & 5930 Switch Series MPLS Configuration Guide-Release 27xx* and *HPE FlexFabric 5940 & 5930 Switch Series MPLS Command Reference-Release 27xx*.

Table 6 MPLS features added in version R2702

| Feature | Command changes |
|--|--|
| Static LSP: Displaying information about static LSPs deployed by OpenFlow | The openflow keyword was added to the display mpls static-lsp command. |
| MPLS L3VPN: Replicating routes from a VPN instance to the public network | The route-replicate from vpn-instance command was added to public instance IPv4 address family view and public instance IPv6 address family view. |
| MPLS L3VPN: Enabling a VPN instance to replicate routes from the public network or other VPN instances | The route-replicate command was added to VPN instance IPv6 address family view. |
| MPLS L3VPN: Creating an OSPF sham | The authentication-none keyword was added to the |

| Feature | Command changes |
|--|---------------------------|
| link with the None authentication mode | sham-link command. |

New features: ACL and QoS features

[Table 7](#) describes the ACL and QoS features added in this software version. For more information about the features and commands, see *HPE FlexFabric 5940 & 5930 Switch Series ACL and QoS Configuration Guide-Release 27xx* and *HPE FlexFabric 5940 & 5930 Switch Series ACL and QoS Command Reference-Release 27xx*.

Table 7 ACL and QoS features added in version R2702

| Feature | Command changes |
|--|--|
| Data buffer: Displaying detailed buffer usage statistics for interfaces | The verbose keyword was added to the display buffer usage interface command. |
| Data buffer: Configuring packet-drop alarms | The following commands were added: <ul style="list-style-type: none"> • buffer packet-drop alarm enable • buffer packet-drop alarm interval |
| Data buffer: Configuring alarm thresholds for the ingress or egress buffer | The following commands were added: <ul style="list-style-type: none"> • buffer threshold alarm enable • buffer threshold alarm interval • buffer usage threshold queue (interface view) • buffer usage threshold queue (system view) |
| Data buffer: Configuring alarm thresholds for the Headroom buffer | The following commands were added: <ul style="list-style-type: none"> • buffer threshold alarm enable • buffer threshold alarm interval • buffer usage threshold headroom • buffer usage threshold headroom queue |

New features: Security features

[Table 8](#) describes the security features added in this software version. For more information about the features and commands, see *HPE FlexFabric 5940 & 5930 Switch Series Security Configuration Guide-Release 27xx* and *HPE FlexFabric 5940 & 5930 Switch Series Security Command Reference-Release 27xx*.

Table 8 Security features added in version R2702

| Feature | Command changes |
|--|--|
| AAA: Setting the version of RADIUS server status change MIB nodes | The radius trap-version command was added. |
| AAA: Specifying a RADIUS server selection mode for reauthentication | The reauthentication server-select command was added. |
| 802.1X: Sending EAP-Success packets on assignment of users to the 802.1X Auth-Fail VLAN or VSI | The dot1x auth-fail eapol command was added. |
| 802.1X: Enabling 802.1X online user | The dot1x server-recovery |

| Feature | Command changes |
|---|--|
| synchronization | online-user-sync command was added. |
| 802.1X: Logging off 802.1X users | The reset dot1x access-user command was added. |
| MAC authentication: Enabling online user synchronization for MAC authentication | The mac-authentication server-recovery online-user-sync command was added. |
| MAC authentication: Logging off MAC authentication users | The reset mac-authentication access-user command was added. |
| Port security: Setting the block timer for blocked MAC addresses | The port-security timer blockmac command was added. |
| SSH: Support for non-interactive method to connect to the SCP server | The no-more-input keyword was added to the scp and scp ipv6 commands. |
| MACsec: Setting the MKA life time | The mka timer mka-life command was added. |

New features: High availability features

[Table 9](#) describes the high availability features added in this software version. For more information about the features and commands, see *HPE FlexFabric 5940 & 5930 Switch Series High Availability Configuration Guide-Release 27xx* and *HPE FlexFabric 5940 & 5930 Switch Series High Availability Command Reference-Release 27xx*.

Table 9 High availability features added in version R2702

| Feature | Command changes |
|------------------------------------|--|
| Track: Configuring a tracked list. | <p>The following commands were added:</p> <ul style="list-style-type: none"> • delay • object • threshold percentage • threshold weight • track list boolean • track list threshold percentage • track list threshold weight |

New features: OpenFlow features

[Table 10](#) describes the OpenFlow features added in this software version. For more information about the features and commands, see *HPE FlexFabric 5940 & 5930 Switch Series OpenFlow Configuration Guide-Release 27xx* and *HPE FlexFabric 5940 & 5930 Switch Series OpenFlow Command Reference-Release 27xx*.

Table 10 OpenFlow features added in version R2702

| Feature | Command changes |
|---|---|
| Configuring an MPLS flow table | The mpls keyword was added to the flow-table command. |
| Reserving ACL resources for extensibility flow tables | The openflow flow-entry |

| Feature | Command changes |
|---------|--|
| | reserve-acl-resource command was added. |

New features: VXLAN features

[Table 11](#) describes the VXLAN features added in this software version. For more information about the features and commands, see *HPE FlexFabric 5940 & 5930 Switch Series VXLAN Configuration Guide-Release 27xx* and *HPE FlexFabric 5940 & 5930 Switch Series VXLAN Command Reference-Release 27xx*.

Table 11 VXLAN features added in version R2702

| Feature | Command changes |
|--|--|
| Enabling BFD on a VXLAN tunnel interface | The tunnel bfd enable command was added. |
| Specifying a global source address for VXLAN tunnels | The tunnel global source-address command was added. |
| Enabling default VXLAN decapsulation | The vxlan default-decapsulation command was added. |

New features: EVPN features

[Table 12](#) describes the EVPN features added in this software version. For more information about the features and commands, see *HPE FlexFabric 5940 & 5930 Switch Series EVPN Configuration Guide-Release 27xx* and *HPE FlexFabric 5940 & 5930 Switch Series EVPN Command Reference-Release 27xx*.

Table 12 EVPN features added in version R2702

| Feature | Command changes |
|--|---|
| EVPN: Enabling ARP mobility event suppression | The evpn route arp-mobility suppression command was added. |
| EVPN: Displaying EVPN ND flood suppression entries | The display evpn route nd suppression command was added. |
| MVXLAN | All MVXLAN commands were newly added. |

Modified feature: Software patching by using issu commands

Feature change description

The number of patch image files supported by **issu** commands for displaying and installing patch images increased from 16 each command to 30 each command.

Command changes

Modified command: display version comp-matrix

Old syntax

```
display version comp-matrix file { boot filename | system filename |  
feature filename&<1-30> | patch filename&<1-16> } *  
  
display version comp-matrix file ipe ipe-filename [ patch  
filename&<1-16> ]
```

New syntax

```
display version comp-matrix file { boot filename | system filename |  
feature filename&<1-30> | patch filename&<1-30> } *  
  
display version comp-matrix file ipe ipe-filename [ patch  
filename&<1-30> ]
```

Views

Any view

Parameters

patch: Specifies a space-separated list of up to 30 patch image files.

filename: Specifies a .bin file in the *filesystemname/filename.bin* format. The file must be stored in the root directory of a file system on the device. The value string excluding the file system location section (if any) can have a maximum of 63 characters. For more information about specifying a file, see file system management in *Fundamentals Configuration Guide*.

Change description

Before modification: The command supports a maximum of 16 patch image files.

After modification: The command supports a maximum of 30 patch image files.

Modified command: issu load

Old syntax

```
issu load file { boot filename | system filename | feature filename&<1-30>  
| patch filename&<1-16> } * slot slot-number&<1-9> [ reboot ]  
  
issu load file ipe ipe-filename [ patch filename&<1-16> ] slot  
slot-number&<1-9> [ reboot ]
```

New syntax

```
issu load file { boot filename | system filename | feature filename&<1-30>  
| patch filename&<1-30> } * slot slot-number&<1-9> [ reboot ]  
  
issu load file ipe ipe-filename [ patch filename&<1-30> ] slot  
slot-number&<1-9> [ reboot ]
```

Views

User view

Parameters

patch: Specifies a space-separated list of up to 30 patch image files.

filename: Specifies a .bin file in the *filesystemname/filename.bin* format. The file must be stored in the root directory of a file system on the device. The value string excluding the file system location

section (if any) can have a maximum of 63 characters. For more information about specifying a file, see file system management in *Fundamentals Configuration Guide*.

Change description

Before modification: The command supports a maximum of 16 patch image files.

After modification: The command supports a maximum of 30 patch image files.

Modified command: issu one-step

Old syntax

```
issu one-step file { boot filename | system filename | feature
filename&<1-30> | patch filename&<1-16> } * [ slot slot-number&<1-9> ]
[ reboot ]
```

```
issu one-step file ipe ipe-filename [ patch filename&<1-16> ] slot
slot-number&<1-9> [ reboot ]
```

New syntax

```
issu one-step file { boot filename | system filename | feature
filename&<1-30> | patch filename&<1-30> } * [ slot slot-number&<1-9> ]
[ reboot ]
```

```
issu one-step file ipe ipe-filename [ patch filename&<1-30> ] slot
slot-number&<1-9> [ reboot ]
```

Views

User view

Parameters

patch: Specifies a space-separated list of up to 30 patch image files.

filename: Specifies a .bin file in the *filesystemname/filename.bin* format. The file must be stored in the root directory of a file system on the device. The value string excluding the file system location section (if any) can have a maximum of 63 characters. For more information about specifying a file, see file system management in *Fundamentals Configuration Guide*.

Change description

Before modification: The command supports a maximum of 16 patch image files.

After modification: The command supports a maximum of 30 patch image files.

Modified feature: Automatic configuration

Feature change description

Before modification: The device supports automatic configuration only on IPv4 networks.

After modification: The device supports automatic configuration on both IPv4 and IPv6 networks.

Command changes

None

Modified feature: Collision handling in BFD MAD

Feature change description

Before modification, BFD MAD uses the following process to handle a multi-active collision:

1. Compares the number of members in each fabric.
2. Sets all fabrics to the Recovery state except the one that has the most members.
3. Compares the member IDs of the masters if all IRF fabrics have the same number of members.
4. Sets all fabrics to the Recovery state except the one that has the lowest numbered master.

After modification, BFD MAD uses the following process to handle a multi-active collision:

1. Compares the health states of split fabrics.
2. Sets all fabrics to the Recovery state except the healthiest one.
3. Compares the number of members in each fabric if all IRF fabrics are in the same health state.
4. Sets all fabrics to the Recovery state except the one that has the most members.
5. Compares the member IDs of their masters if all IRF fabrics have the same number of members.
6. Sets all fabrics to the Recovery state except the one that has the lowest numbered master.

Command changes

None

Modified feature: Processing after the link mode of an Ethernet interface is switched

Feature change description

Before modification: After the link mode of an Ethernet interface is switched, all commands except the **shutdown** command are restored to their defaults.

After modification: After the link mode of an Ethernet interface is switched, all commands except the **shutdown** and **speed** command are restored to their defaults.

Command changes

None.

Modified feature: Link state change suppression on an interface

Feature change description

In this version, the syntax for configuring link state change suppression on an Ethernet interface is modified. When the link state change suppression interval is configured in seconds on an Ethernet interface, the value range is modified.

Command changes

Modified command: link-delay

Old syntax

```
link-delay [ msec ] delay-time [ mode { up | updown } ]  
undo link-delay [ msec ] delay-time [ mode { up | updown } ]
```

New syntax

```
link-delay { down | up } [ msec ] delay-time  
undo link-delay { down | up }
```

Views

Ethernet interface view

Change description

Before modification:

- If the **mode** keyword is not specified, the link-down events are suppressed.
- If the **mode up** keyword combination is specified, the link-up events are suppressed.
- If the **mode updown** keyword combination is specified, both link-down and link-up events are suppressed.
- If the suppression interval configured in the command without the **mode** keyword specified is the same as the suppression interval configured in the command with the **mode up** keyword combination specified on an interface, the two commands are automatically merged into the command with the **mode updown** keyword combination specified in the configuration file of the interface.
- If the **msec** keyword is not specified, the link state change suppression interval is configured in seconds, and the value range is 0 to 30.

After modification:

- If the **down** keyword is not specified, the link-down events are suppressed.
- If the **up** keyword is specified, the link-up events are suppressed.
- You can set different link state change suppression intervals for link-down events and link-up events.
- If the **msec** keyword is not specified, the link state change suppression interval is configured in seconds, and the value range is 0 to 300.

Modified feature: Setting the global aging timer for ND entries in stale state

Feature change description

This release supports setting the global aging timer in seconds for ND entries in stale state.

Command changes

Modified command: `ipv6 neighbor stale-aging`

Old syntax

```
ipv6 neighbor stale-aging aging-time
```

New syntax

```
ipv6 neighbor stale-aging { aging-minutes | second aging-seconds }
```

Views

System view

Change description

Before modification: The global aging timer can be set in minutes for ND entries in stale state.

After modification: The global aging timer can be set in minutes or in seconds for ND entries in stale state.

Modified feature: Setting the interface-specific aging timer for ND entries in stale state

Feature change description

This release supports setting the interface-specific aging timer in seconds for ND entries in stale state.

Command changes

Modified command: `ipv6 neighbor timer stale-aging`

Old syntax

```
ipv6 neighbor timer stale-aging aging-time
```

New syntax

```
ipv6 neighbor timer stale-aging { aging-minutes | second aging-seconds }
```

Views

Layer 3 Ethernet interface/subinterface view

Layer 3 aggregate interface/subinterface view

VXLAN VSI interface view

VLAN interface view

Tunnel interface view

Change description

Before modification: The interface-specific aging timer can be set in minutes for ND entries in stale state.

After modification: The interface-specific aging timer can be set in minutes or in seconds for ND entries in stale state.

Modified feature: Configuring OSPF area authentication

Feature change description

From this release, a key ID becomes optional in HMAC-MD5 and MD5 authentication modes, and a key string becomes optional in simple authentication mode.

Command changes

Modified command: authentication-mode

Old syntax

In MD5/HMAC-MD5 authentication mode:

```
authentication-mode { hmac-md5 | md5 } key-id { cipher | plain } string
undo authentication-mode [ { hmac-md5 | md5 } key-id ]
```

In simple authentication mode:

```
authentication-mode simple { cipher | plain } string
undo authentication-mode
```

New syntax

In MD5/HMAC-MD5 authentication mode:

```
authentication-mode { hmac-md5 | md5 } [ key-id { cipher | plain } string ]
undo authentication-mode [ { hmac-md5 | md5 } [ key-id ] ]
```

In simple authentication mode:

```
authentication-mode simple [ { cipher | plain } string ]
undo authentication-mode
```

Views

OSPF area view

Change description

Before modification: A key ID is required in HMAC-MD5 and MD5 authentication modes, and a key string is required in simple authentication mode.

After modification: A key ID is optional in HMAC-MD5 and MD5 authentication modes, and a key string is optional in simple authentication mode.

Modified feature: Configuring OSPF interface authentication

Feature change description

From this release, a key ID becomes optional in HMAC-MD5 and MD5 authentication modes, a key string becomes optional in simple authentication mode, and the none authentication mode is supported.

Command changes

Modified command: ospf authentication-mode

Old syntax

In MD5/HMAC-MD5 authentication mode:

```
ospf authentication-mode { hmac-md5 | md5 } key-id { cipher | plain } string
undo ospf authentication-mode { hmac-md5 | md5 } key-id
```

In simple authentication mode:

```
ospf authentication-mode simple { cipher | plain } string
undo ospf authentication-mode simple
```

New syntax

In MD5/HMAC-MD5 authentication mode:

```
ospf authentication-mode { hmac-md5 | md5 } [ key-id { cipher | plain }
string ]
undo ospf authentication-mode { hmac-md5 | md5 } [ key-id ]
```

In simple authentication mode:

```
ospf authentication-mode simple [ { cipher | plain } string ]
undo ospf authentication-mode simple
```

In none authentication mode:

```
ospf authentication-mode none
undo ospf authentication-mode none
```

Views

Interface view

Change description

Before modification: A key ID is required in HMAC-MD5 and MD5 authentication modes, a key string is required in simple authentication mode, and the none authentication mode is not supported.

After modification: A key ID is optional in HMAC-MD5 and MD5 authentication modes, a key string is optional in simple authentication mode, and the none authentication mode is supported.

Modified feature: Configuring a virtual link

Feature change description

From this release, a key ID becomes optional in HMAC-MD5 and MD5 authentication modes, a key string becomes optional in simple authentication mode, and the none authentication mode is supported.

Command changes

Modified command: vlink-peer

Old syntax

```
vlink-peer router-id [ dead seconds | hello seconds | { { hmac-md5 | md5 }  
key-id { cipher | plain } string | keychain keychain-name | simple { cipher |  
plain } string } | retransmit seconds | trans-delay seconds ] *  
  
undo vlink-peer router-id [ dead | hello | { hmac-md5 | md5 } key-id |  
keychain | retransmit | simple | trans-delay ] *
```

New syntax

```
vlink-peer router-id [ dead seconds | hello seconds | [ authentication-none  
/ { hmac-md5 | md5 } [ key-id { cipher | plain } string ] | keychain  
keychain-name | simple [ { cipher | plain } string ] ] | retransmit seconds  
| trans-delay seconds ] *  
  
undo vlink-peer router-id [ dead | hello | [ authentication-none /  
{ hmac-md5 | md5 } [ key-id ] | keychain ] | retransmit | simple | trans-delay ]  
*
```

Views

Interface view

Change description

Before modification: A key ID is required in HMAC-MD5 and MD5 authentication modes, a key string is required in simple authentication mode, and the none authentication mode is not supported.

After modification: A key ID is optional in HMAC-MD5 and MD5 authentication modes, a key string is optional in simple authentication mode, and the none authentication mode is supported.

Modified feature: Displaying statistics for the IGMP messages and PIMv2 hello messages learned through IGMP snooping

Feature change description

The command to display statistics for the IGMP messages and PIMv2 hello messages learned through IGMP snooping has the following changes:

- The layout of the command output was changed.
- The **DRNI** field indicating the number of IGMP messages received on the DRNI IPP was added to the command output.

Command changes

Modified command: display igmp-snooping statistics

Syntax

```
display igmp-snooping statistics
```

Views

Any view

Change description

The **DRNI** field was added to the command output. This field indicates the number of IGMP messages received on the DRNI IPP.

Before modification:

Display statistics for the IGMP messages and PIMv2 hello messages learned through IGMP snooping.

```
<Sysname> display igmp-snooping statistics
Received IGMP general queries: 0
Received IGMPv1 reports: 0
Received IGMPv2 reports: 19
Received IGMP leaves: 0
Received IGMPv2 specific queries: 0
Sent IGMPv2 specific queries: 0
Received IGMPv3 reports: 1
Received IGMPv3 reports with right and wrong records: 0
Received IGMPv3 specific queries: 0
Received IGMPv3 specific sg queries: 0
Sent IGMPv3 specific queries: 0
Sent IGMPv3 specific sg queries: 0
Received PIMv2 hello: 0
Received error IGMP messages: 19
```

Table 13 Command output

| Field | Description |
|--------------------------------------|--|
| general queries | Number of IGMP general queries. |
| specific queries | Number of IGMP group-specific queries. |
| reports | Number of IGMP reports. |
| leaves | Number of IGMP leave messages. |
| reports with right and wrong records | Number of IGMP reports with correct and incorrect records. |
| specific sg queries | Number of IGMP group-and-source-specific queries. |
| PIMv2 hello | Number of PIMv2 hello messages. |
| error IGMP messages | Number of error IGMP messages. |

After modification:

Display statistics for the IGMP messages and PIMv2 hello messages learned through IGMP snooping.

```
<Sysname> display igmp-snooping statistics
Received IGMP messages:
  IGMP general queries: 19
  DRNI: 19
  IGMPv2 specific queries: 10
  DRNI: 0
  IGMPv3 specific queries: 0
```

```

DRNI: 0
IGMPv3 specific sg queries: 0
DRNI: 0
IGMPv1 reports: 0
DRNI: 0
IGMPv2 reports: 19
DRNI: 19
IGMPv3 reports: 0
DRNI: 0
IGMPv3 reports with right and wrong records: 0
DRNI: 0
IGMP leaves: 5
DRNI: 5
Error IGMP messages: 1
DRNI: 1
Sent IGMP messages:
IGMPv2 specific queries: 0
IGMPv3 specific queries: 0
IGMPv3 specific sg queries: 0

Received PIMv2 hello : 0

```

Table 14 Command output

| Field | Description |
|--------------------------------------|--|
| general queries | Number of IGMP general queries. |
| specific queries | Number of IGMP group-specific queries. |
| reports | Number of IGMP reports. |
| leaves | Number of IGMP leave messages. |
| reports with right and wrong records | Number of IGMP reports with correct and incorrect records. |
| specific sg queries | Number of IGMP group-and-source-specific queries. |
| error IGMP messages | Number of IGMP messages with errors. |
| DRNI | IGMP messages received on the DRNI IPP. |

Modified feature: Displaying statistics for the MLD messages and IPv6 PIM hello messages learned through MLD snooping

Feature change description

The command to display statistics for the MLD messages and IPv6 PIM hello messages learned through MLD snooping has the following changes:

- The layout of the command output was changed.
- The **DRNI** field indicating the number of MLD messages received on the DRNI IPP was added to the command output.

Command changes

Modified command: display mld-snooping statistics

Syntax

display mld-snooping statistics

Views

Any view

Change description

The **DRNI** field was added to the command output. This field indicates the number of MLD messages received on the DRNI IPP.

Before modification:

Display statistics for the MLD messages and IPv6 PIM hello messages learned through MLD snooping.

```
<Sysname> display mld-snooping statistics
Received MLD general queries:  0
Received MLDv1 specific queries:  0
Received MLDv1 reports:  0
Received MLD done:  0
Sent      MLDv1 specific queries:  0
Received MLDv2 reports:  0
Received MLDv2 reports with right and wrong records:  0
Received MLDv2 specific queries:  0
Received MLDv2 specific sg queries:  0
Sent      MLDv2 specific queries:  0
Sent      MLDv2 specific sg queries:  0
Received IPv6 PIM hello:  0
Received error MLD messages:  0
```

Table 15 Command output

| Field | Description |
|--------------------------------------|--|
| general queries | Number of MLD general queries. |
| specific queries | Number of MLD multicast-address-specific queries. |
| reports | Number of MLD reports. |
| done | Number of MLD done messages. |
| reports with right and wrong records | Number of MLD reports with correct and incorrect records. |
| specific sg queries | Number of MLD multicast-address-and-source-specific queries. |
| IPv6 PIM hello | Number of IPv6 PIM hello messages. |
| error MLD messages | Number of MLD messages with errors. |

After modification:

Display statistics for the MLD messages and IPv6 PIM hello messages learned through MLD snooping.

```
<Sysname> display mld-snooping statistics
```

```

Received MLD messages:
  MLD general queries: 19
    DRNI: 19
  MLDv1 specific queries: 0
    DRNI: 0
  MLDv2 specific queries: 0
    DRNI: 0
  MLDv2 specific sg queries: 0
    DRNI: 0
  MLDv1 reports: 19
    DRNI: 19
  MLDv2 reports: 0
    DRNI: 0
  MLDv2 reports with right and wrong records: 0
    DRNI: 0
  MLD dones: 5
    DRNI: 5
  Error MLD messages: 1
    DRNI: 1
Sent MLD messages:
  MLDv1 specific queries: 0
  MLDv2 specific queries: 0
  MLDv2 specific sg queries: 0

```

Received IPv6 PIM hello: 0

Table 16 Command output

| Field | Description |
|--------------------------------------|--|
| general queries | Number of MLD general queries. |
| specific queries | Number of MLD multicast-address-specific queries. |
| reports | Number of MLD reports. |
| dones | Number of MLD done messages. |
| reports with right and wrong records | Number of MLD reports with correct and incorrect records. |
| specific sg queries | Number of MLD multicast-address-and-source-specific queries. |
| error MLD messages | Number of MLD messages with errors. |
| DRNI | MLD messages received on the DRNI IPP. |

Modified feature: Enabling a VPN instance to replicate routes from the public network or other VPN instances

Feature change description

This release allows a VPN instance to advertise the BGP routes replicated from the public network or another VPN instance.

Command changes

Modified command: route-replicate

Old syntax

```
route-replicate from { public | vpn-instance vpn-instance-name } protocol  
bgp as-number [ route-policy route-policy-name ]
```

New syntax

```
route-replicate from { public | vpn-instance vpn-instance-name } protocol  
bgp as-number [ advertise ] [ route-policy route-policy-name ]
```

Views

VPN instance IPv4 address family view

Change description

Before modification: The **advertise** keyword is not available when you enable a VPN instance to replicate BGP routes from the public network or another VPN instance.

After modification: The **advertise** keyword is available when you enable a VPN instance to replicate BGP routes from the public network or another VPN instance. This keyword allows the VPN instance to advertise the replicated BGP routes.

Modified feature: Using BFD to verify LSP or MPLS TE tunnel connectivity

Feature change description

The value range for the local discriminator of the BFD session was changed.

Command changes

Modified command: mpls bfd

Syntax

```
mpls bfd dest-addr mask-length [ nexthop nexthop-address [ discriminator  
local local-id remote remote-id ] ] [ template template-name ]  
mpls bfd [ [ discriminator local local-id remote remote-id ] [ template  
template-name ] | echo ]
```

Views

System view

Tunnel interface view

Change description

Before modification: The value range for the *local-id* argument is 1 to 64.

After modification: The value range for the *local-id* argument is 1 to 500.

Modified feature: MAC authentication VLAN mode

Feature change description

As from this version, the device changed its manner in handling an online authenticated user that meets the following requirements:

- The user accesses a MAC authentication-enabled port in single-VLAN mode from a VLAN different from the VLAN where it was authenticated the last time.
- The user has been assigned an authorization VLAN.

Command changes

Modified command: mac-authentication host-mode multi-vlan

Syntax

```
mac-authentication host-mode multi-vlan
```

Views

Layer 2 Ethernet interface view

Change description

The device changed its manner in handling an online authenticated user that meets the following requirements:

- The user accesses a MAC authentication-enabled port in single-VLAN mode from a VLAN different from the VLAN where it was authenticated the last time.
- The user has been assigned an authorization VLAN.

Before modification: When the user accesses the port from a new VLAN, the device does not authenticate the user or log off the user from the authorization VLAN. The user is online in the authorization VLAN.

After modification: When the user accesses the port from a new VLAN, the device handles the user depending on the status of the port security MAC move feature.

- If port security MAC move is disabled, the user cannot pass authentication and come online from the new VLAN until after it goes offline from the port.
- If port security MAC move is enabled, the user can pass authentication on the new VLAN and come online without having to first go offline from the port. After the user passes authentication on the new VLAN, the original authentication session of the user is deleted from the port.

No changes were introduced to the command syntax.

NOTE:

To enable single-VLAN mode, execute the **undo** form of the **mac-authentication host-mode multi-vlan** command.

To enable the port security MAC move feature, use the **port-security mac-move permit** command.

Modified feature: Port security MAC move

Feature change description

As from this version, the port security MAC move feature takes effect on users that move between VLANs on a port in addition to users that move between ports.

Command changes

Modified command: port-security mac-move permit

Syntax

```
port-security mac-move permit
```

Views

System view

Change description

Before modification: Port security MAC move setting takes effect only on users that move between ports on the device.

After modification: Port security MAC move setting also take effect on users that move between VLANs on a port in a VXLAN environment. In additional, user packets are VLAN tagged.

- If this feature is disabled, authenticated users must go offline from the original VLAN first before they can be reauthenticated successfully on the new VLAN and come online.
- If this feature is enabled, authenticated users can be reauthenticated successfully on the new VLAN without having to go offline from the original VLAN. The port will remove the users from the original VLAN immediately after the users are reauthenticated successfully on the new VLAN.

No changes were introduced to the command syntax.

NOTE:

MAC authentication multi-VLAN mode has higher priority than MAC move for users moving between VLANs on a port. If MAC authentication multi-VLAN mode is enabled, these users can come online in the new VLAN without being reauthenticated. To enable MAC authentication multi-VLAN mode, use the `mac-authentication host-mode multi-vlan` command.

Modified feature: Creating a BFD session for detecting the local interface state

Feature change description

From this release, the value range for the local discriminator of the BFD session is changed.

Command changes

Modified command: bfd detect-interface

Syntax

```
bfd detect-interface source-ip ip-address [ discriminator local  
local-value remote remote-value ] [ template template-name ]
```

Views

Interface view

Change description

Before modification: The value range for the *local-value* argument is 193 to 256.

After modification: The value range for the *local-value* argument is 1501 to 2000.

Modified feature: Associating Track with application modules

Feature change description

From this release, you cannot configure the notification delay when associating Track with application modules. Creating a track entry associated with an application module enters Track view. You can configure the delay only in Track view for notifying the application module of track entry state changes.

Command changes

Modified command: track bfd ctrl

Old syntax

```
track track-entry-number bfd ctrl [ interface interface-type  
interface-number | vpn-instance vpn-instance-name ] remote ip  
remote-ip-address local ip local-ip-address [ delay { negative  
negative-time | positive positive-time } * ]
```

New syntax

```
track track-entry-number bfd ctrl [ interface interface-type  
interface-number | vpn-instance vpn-instance-name ] remote ip  
remote-ip-address local ip local-ip-address
```

Change description

Before modification: This command allows you to set the notification delay when associating Track with the application module, and does not enter Track view.

After modification: Creating a track entry associated with the application module enters Track view. You can configure the notification delay in Track view.

Modified command: track bfd echo

Old syntax

```
track track-entry-number bfd echo interface interface-type  
interface-number remote ip remote-ip-address local ip local-ip-address  
[ delay { negative negative-time | positive positive-time } * ]
```

New syntax

```
track track-entry-number bfd echo interface interface-type  
interface-number remote ip remote-ip-address local ip local-ip-address
```

Change description

Before modification: This command allows you to set the notification delay when associating Track with the application module, and does not enter Track view.

After modification: Creating a track entry associated with the application module enters Track view. You can configure the notification delay in Track view.

Modified command: track cfd

Old syntax

```
track track-entry-number cfd cc service-instance instance-id mep mep-id  
[ delay { negative negative-time | positive positive-time } * ]
```

New syntax

```
track track-entry-number cfd cc service-instance instance-id mep mep-id
```

Change description

Before modification: This command allows you to set the notification delay when associating Track with the application module, and does not enter Track view.

After modification: Creating a track entry associated with the application module enters Track view. You can configure the notification delay in Track view.

Modified command: track interface

Old syntax

```
track track-entry-number interface interface-type interface-number  
[ delay { negative negative-time | positive positive-time } * ]
```

New syntax

```
track track-entry-number interface interface-type interface-number
```

Change description

Before modification: This command allows you to set the notification delay when associating Track with the application module, and does not enter Track view.

After modification: Creating a track entry associated with the application module enters Track view. You can configure the notification delay in Track view.

Modified command: track interface physical

Old syntax

```
track track-entry-number interface interface-type interface-number  
physical [ delay { negative negative-time | positive positive-time } * ]
```

New syntax

```
track track-entry-number interface interface-type interface-number  
physical
```

Change description

Before modification: This command allows you to set the notification delay when associating Track with the application module, and does not enter Track view.

After modification: Creating a track entry associated with the application module enters Track view. You can configure the notification delay in Track view.

Modified command: track interface protocol

Old syntax

```
track track-entry-number interface interface-type interface-number  
protocol { ipv4 | ipv6 } [ delay { negative negative-time | positive  
positive-time } * ]
```

New syntax

```
track track-entry-number interface interface-type interface-number  
protocol { ipv4 | ipv6 }
```

Change description

Before modification: This command allows you to set the notification delay when associating Track with the application module, and does not enter Track view.

After modification: Creating a track entry associated with the application module enters Track view. You can configure the notification delay in Track view.

Modified command: track ip route reachability

Old syntax

```
track track-entry-number ip route [ vpn-instance vpn-instance-name ]  
ip-address { mask-length | mask } reachability [ delay { negative  
negative-time | positive positive-time } * ]
```

New syntax

```
track track-entry-number ip route [ vpn-instance vpn-instance-name ]  
ip-address { mask-length | mask } reachability
```

Change description

Before modification: This command allows you to set the notification delay when associating Track with the application module, and does not enter Track view.

After modification: Creating a track entry associated with the application module enters Track view. You can configure the notification delay in Track view.

Modified command: track lldp neighbor

Old syntax

```
track track-entry-number lldp neighbor interface interface-type  
interface-number [ delay { negative negative-time | positive  
positive-time } * ]
```

New syntax

```
track track-entry-number lldp neighbor interface interface-type  
interface-number
```

Change description

Before modification: This command allows you to set the notification delay when associating Track with the application module, and does not enter Track view.

After modification: Creating a track entry associated with the application module enters Track view. You can configure the notification delay in Track view.

Modified command: track nqa

Old syntax

```
track track-entry-number nqa entry admin-name operation-tag reaction  
item-number [ delay { negative negative-time | positive positive-time } * ]
```

New syntax

```
track track-entry-number nqa entry admin-name operation-tag reaction  
item-number
```

Change description

Before modification: This command allows you to set the notification delay when associating Track with the application module, and does not enter Track view.

After modification: Creating a track entry associated with the application module enters Track view. You can configure the notification delay in Track view.

Modified feature: Configuring an interface event for a CLI-defined EAA monitor policy

Feature change description

This release added support of monitoring traffic statistics on multiple interfaces. The interfaces in an interface range must be the same type.

Command changes

Modified command: event-interface

Old syntax

```
event interface interface-type interface-number monitor-obj monitor-obj  
start-op start-op start-val start-val restart-op restart-op restart-val  
restart-val [ interval interval ]
```

New syntax

```
event interface interface-list monitor-obj monitor-obj start-op start-op  
start-val start-val restart-op restart-op restart-val restart-val  
[ interval interval ]
```

Change description

Before modification: Only one interface can be monitored.

After modification: Multiple interfaces can be monitored. The interfaces in an interface range must be the same type. The start interface number must be smaller than the end interface number.

Modified feature: Configuring an EAA monitor policy by using Tcl

Feature change description

In this release, the system uses the platformtools module rather than the Comware module to create a Tcl script file.

Command changes

N/A

Modified feature: Removing a TCP or UDP listening service for a VPN instance

Feature change description

From this release, you can remove a TCP or UDP listening service for the specified VPN instance.

Command changes

Modified command: nqa server tcp-connect

Old syntax

```
nqa server tcp-connect ip-address port-number [ vpn-instance  
vpn-instance-name ] [ tos tos ]  
  
undo nqa server tcp-connect ip-address port-number
```

New syntax

```
nqa server tcp-connect ip-address port-number [ vpn-instance  
vpn-instance-name ] [ tos tos ]  
  
undo nqa server tcp-connect ip-address port-number [ vpn-instance  
vpn-instance-name ]
```

Views

System view

Change description

The **vpn-instance** *vpn-instance-name* option was added to the **undo** form of the command.

vpn-instance *vpn-instance-name*: Specifies an MPLS L3VPN instance by its name, a case-sensitive string of 1 to 31 characters. If you do not specify this option, the NQA server listens to the TCP services on an IP address on the public network.

Modified command: nqa server udp-echo

Old syntax

```
nqa server udp-echo ip-address port-number [ vpn-instance  
vpn-instance-name ] [ tos tos ]  
  
undo nqa server udp-echo ip-address port-number
```

New syntax

```
nqa server udp-echo ip-address port-number [ vpn-instance  
vpn-instance-name ] [ tos tos ]  
  
undo nqa server udp-echo ip-address port-number [ vpn-instance  
vpn-instance-name ]
```

Views

System view

Change description

The **vpn-instance** *vpn-instance-name* option was added to the **undo** form of the command.

vpn-instance *vpn-instance-name*: Specifies an MPLS L3VPN instance by its name, a case-sensitive string of 1 to 31 characters. If you do not specify this option, the NQA server listens to UDP services on an IP address on the public network.

Modified feature: Specifying the role of the device in the VCF fabric

Feature change description

In this version and later, a data center VCF fabric supports the spine-aggregate-leaf topology. Residing on the distribution layer, an aggregate node is between leaf nodes and spine nodes. In such a VCF fabric, OSPF runs on the Layer 3 networks between the spine and aggregate nodes and between the aggregate and leaf nodes. The default role of the device in the VCF fabric is aggregate.

Automated overlay network deployment is not supported on aggregate nodes.

Command reference

Modified command: vcf-fabric role

Old syntax

```
vcf-fabric role { access | leaf | spine }  
undo vcf-fabric role
```

New syntax

```
vcf-fabric role { access | aggr | leaf | spine }  
undo vcf-fabric role
```

Views

System view

Change description

The **aggr** keyword was added.

aggr: Specifies the aggregate node.

Modified command: display vcf-fabric role

Syntax

```
display vcf-fabric role
```

Views

Any view

Change description

The **aggr** value was added to the **Default role** and **Current role** fields.

Before modification:

Display the role of the device in the VCF fabric.

```
<System> display vcf-fabric role
```

Default role: leaf

Current role: leaf

After modification:

Display the role of the device in the VCF fabric.

```
<System> display vcf-fabric role
```

Default role: leaf

Current role: **aggr**

Table 17 Command output

| Field | description |
|--------------|--|
| Default role | <p>Default role of the device in the VCF fabric:</p> <ul style="list-style-type: none">• Not specified—The device does not have a role.• spine—The device is a spine node.• aggr—The device is an aggregate node.• leaf—The device is a leaf node.• access—The device is an access node. <p>NOTE:</p> <p>The aggr field was added in this version.</p> |
| Current role | <p>Current role of the device in the VCF fabric:</p> <ul style="list-style-type: none">• Not specified—The device does not have a role.• spine—The device is a spine node.• aggr—The device is an aggregate node.• leaf—The device is a leaf node.• access—The device is an access node. <p>NOTE:</p> <p>The aggr field was added in this version.</p> |

Modified command: display vcf-fabric underlay autoconfigure

Syntax

display vcf-fabric underlay autoconfigure

Views

Any view

Change description

Before modification: VCF fabric does not support the spine-aggregate-leaf topology.

After modification: VCF fabric supports the spine-aggregate-leaf topology. If you use this command on a master spine node, BGP peer information on all spine and leaf nodes will be displayed.

Display information about automated underlay network deployment.

```
<Sysname> display vcf-fabric underlay autoconfigure
```

success command:

```
#
system
clock timezone beijing add 08:00:00
#
system
lldp global enable
lldp compliance cdp
#
system
ospf 1
graceful-restart ietf
area 0.0.0.0
#
system
interface LoopBack0
#
system
l2vpn enable
#
system
vxlan tunnel mac-learning disable
vxlan tunnel arp-learning disable
#
system
stp global enable
#
system
ntp-service enable
ntp-service unicast-peer 172.16.1.136
#
system
netconf soap https enable
restful https enable
#
```

```

system
info-center loghost 172.16.1.136
#
system
local-user aaa
password *****
service-type https
authorization-attribute user-role network-admin
#
system
line vty 0 63
authentication-mode scheme
user-role network-admin
#
system
bgp 100
graceful-restart
address-family l2vpn evpn
#
system
vcf-fabric topology enable
#
system
neutron
rabbit user openstack
rabbit password *****
rabbit host ip 172.16.1.136
restful user aaa password *****
vpn-target 1:1 export-extcommunity
vsi-mac 789c-2f5f-0200
network-type distributed-vxlan
proxy-arp enable
l2agent enable
l3agent enable
#
system
snmp-agent
snmp-agent packet max-size 4096
snmp-agent community read public
snmp-agent community write private
snmp-agent sys-info version all
#
Uplink interface:
Ten-GigabitEthernet1/0/1
Ten-GigabitEthernet1/0/2
Downlink interface:
Ten-GigabitEthernet1/0/3
Ten-GigabitEthernet1/0/4

```

```

Loopback0 IP allocation:
    Device_MAC      Loopback_IP      Management_IP      State
    a43c-adae-0400  200.1.1.6      172.16.1.193      Up
    a43c-9aa7-0100  200.1.1.4      172.16.1.152      Up
Loopback1 IP allocation:
    Device_MAC      Loopback_IP      Management_IP      State
    a43c-adae-0400  100.1.1.6      172.16.1.193      Up
    a43c-9aa7-0100  100.1.1.4      172.16.1.152      Up
Aggregation configuration:
    AggrID: 7
    interface: Ten-GigabitEthernet1/0/5
    interface: Ten-GigabitEthernet1/0/6
IRF allocation:
    Self Bridge Mac: a43c-adae-0400
    IRF Status: Yes
    Member List: [1, 2]
BGP peer configuration:
    Device_MAC: 265b-208d-0300
    Device_role:leaf
    BGP peer list:
        200.1.4.5
    Device_MAC: 1ea6-06a4-0500
    Device_role:spine
    BGP peer list:
        200.1.4.4
        200.1.4.7
    Device_MAC: lead-aed1-0600
    Device_role:leaf
    BGP peer list:
        200.1.4.5

```

Table 18 Command output

| Field | Description |
|----------------------------|--|
| success command | Commands that have been successfully executed during automated underlay network deployment. |
| Uplink interface | Uplink interfaces of the device. |
| Downlink interface | Downlink interfaces of the device. |
| Loopback0 IP allocation | IP addresses assigned to Loopback 0. |
| Loopback1 IP allocation | IP addresses assigned to Loopback 1. |
| Aggregation configurations | Information about Layer 2 aggregation groups. |
| IRF allocation | IRF configurations, including IRF bridge MAC address of the IRF fabric, IRF status, and the IRF member ID of the device. |
| BGP peer configuration | Information about BGP peers of the device. NOTE: In this version and later, BGP peer information on all spine and leaf nodes is displayed only when this command is executed on the master spine node in a spine-aggregate-leaf VCF fabric. |

Modified feature: Frame match criteria of VXLAN Ethernet service instances

Feature change description

The frame match criteria were changed for Ethernet service instances in VLAN access mode.

Command changes

Modified command: encapsulation

Syntax

```
encapsulation s-vid vlan-id [ c-vid { vlan-id-list | all } | only-tagged ]
encapsulation s-vid vlan-id-list [ c-vid vlan-id-list ]
encapsulation { default | tagged | untagged }
undo encapsulation
```

Views

Ethernet service instance view

Change description

Before modification: On a Layer 2 Ethernet or aggregate interface, if you configure the **encapsulation default** criterion for an Ethernet service instance in VLAN access mode, make sure no other Ethernet service instances exist on the interface.

After modification: On a Layer 2 Ethernet or aggregate interface, if an Ethernet service instance in VLAN access mode uses the **encapsulation default** criterion, you can create other Ethernet service instances on the interface. The interface matches traffic to the Ethernet service instances in the following criterion order:

- **encapsulation { tagged | untagged }**
- **encapsulation default**
- **encapsulation s-vid vlan-id [c-vid { vlan-id | all } | only-tagged]** and **encapsulation s-vid vlan-id-list [c-vid vlan-id-list]**

Modified feature: NETCONF logging

Feature change description

The default setting for NETCONF logging was changed.

Before modification: NETCONF logging is disabled.

After modification: NETCONF logging is enabled only for row operations for <action> and <edit-config> operations.

Command changes

Modified command: netconf log

Old syntax

```
netconf log source { all | { agent | soap } * } { protocol-operation { all  
| { action | config | get | set | session | syntax | others } * } |  
row-operation | verbose }  
  
undo netconf log source { all | { agent | soap } * } { protocol-operation  
{ all | { action | config | get | session | set | syntax | others } * }  
| row-operation | verbose }
```

New syntax

```
netconf log source { all | { agent | soap } * } { protocol-operation { all  
| { action | config | get | set | session | syntax | others } * } | verbose }  
  
undo netconf log source { all | { agent | soap } * } { protocol-operation  
{ all | { action | config | get | session | set | syntax | others } * }  
| verbose }
```

Views

System view

Change description

Before modification: NETCONF logging is disabled by default. To enable the device to log row operations for <action> and <edit-config> operations, you must execute the **netconf log source { all | { agent | soap } * } protocol-operation** command.

After modification: NETCONF logging is enabled only for row operations of <action> and <edit-config> operations by default. You do not need to execute the **netconf log source { all | { agent | soap } * } protocol-operation** command.

Feature 2701

This release has the following changes:

- New feature: Filtering the output from a display command by using multiple conditions
- New feature: Executing a batch file
- New feature: Disabling automatic system-wide next-startup configuration file operations
- New feature: Using an encrypted configuration file to roll back configuration
- New feature: Support for encrypted configuration files for configuration comparison
- New feature: ND MAD on management Ethernet ports
- New feature: Fast retrain
- New feature: Displaying the status and packet statistics of interfaces
- New feature: Displaying operating status and information of all interfaces except subinterfaces
- New feature: Enabling PFC on all interfaces
- New feature: Configuring automatic link aggregation
- New feature: Disabling the default action of selecting a Selected port for dynamic aggregation groups that have not received LACPDUs
- New feature: Physical state change suppression on aggregate interfaces
- New feature: VXLAN multiport ARP entries
- New feature: Enabling recording user port move for ARP
- New feature: Enabling recording user IP address conflicts
- New feature: Enabling interface consistency check for ARP and MAC address entries
- New feature: Specifying a DHCP relay address pool for DHCP clients
- New feature: Configuring the padding mode for the Vendor-Specific sub-option
- New feature: Configuring IPv6 ND direct route advertisement
- New feature: Configuring ND snooping in a VXLAN
- New feature: Enabling recording user IPv6 address conflicts
- New feature: Enabling recording user port move for ND
- New feature: Enabling ND logging for user online and offline events
- New feature: Specifying the boot file URL in RA messages
- New feature: Setting the aging timer for ND entries in stale state on an interface
- New feature: Specifying DNS server information in RA messages
- New feature: Specifying the DNS domain name suffix information in RA messages
- New feature: Suppressing DNS information in RA messages
- New feature: Enabling the DHCPv6 relay agent to advertise host routes for assigned IPv6 addresses
- New feature: Enabling IPv6 address binding conversion for IP source guard
- New feature: Configuring DHCPv6 guard
- New feature: Ignoring router IDs during optimal route selection
- New feature: Redistributing the local network specified in the public instance or a VPN instance
- New feature: Apply route update interval setting to withdrawn routes
- New feature: Enabling prioritized withdrawal of specific routes

- New feature: Configuring BGP RPKI
- New feature: Configuring BGP policy accounting
- New feature: Configuring BGP to immediately send updates for routes that match an IPv6 prefix list
- New feature: BGP optimal route selection delay
- New feature: Specifying a VPN instance for the BMP server
- New feature: Routing policy change delay
- New feature: Support of a routing policy for L3VNI configuration
- New feature: Configuring the BGP RPKI validation state match criterion for a routing policy node
- New feature: Enabling logging for source MAC-based ARP attack detection
- New feature: Enabling ND attack detection logging
- New feature: Configuring BFD session flapping suppression
- New feature: Configuring the timer that delays reporting the first BFD session establishment failure to the data link layer
- New feature: Enabling special processing for BFD sessions
- New feature: Sbfd
- New feature: Associating Track with a control-mode BFD session
- New feature: Enabling the device to send IGMP general queries upon a path change
- New feature: Enabling the device to send MLD general queries upon a path change
- New feature: Enabling embedded RP
- New feature: Specifying a local network to be advertised in the public instance or a VPN instance
- New feature: Changing the next hop address of VPNv4 routes to a VPN address
- New feature: Displaying tunnel policy information
- New feature: Sbfd for MPLS
- New feature: MPLS SR
- New feature: Displaying buffer usage statistics for interfaces
- New feature: Configuring an EAP profile
- New feature: AAA test
- New feature: Displaying and clearing history authentication and accounting load statistics for all RADIUS servers
- New feature: 802.1X offline detection
- New feature: 802.1X unauthenticated user aging
- New feature: 802.1X EAP-TLS fragmentation for packets sent to the server
- New feature: User aging for unauthenticated MAC authentication users
- New feature: MAC authentication method
- New feature: User-specific MAC authentication offline detection
- New feature: NAS-Port-Type attribute
- New feature: Checking the issuing of category-2 portal filtering rules
- New feature: Configuring user authentication timeout timer for global password control
- New feature: Configuring password control for local network access users
- New feature: Configuring password change at first login

- New feature: Enabling the SSL server to send the complete certificate chain during SSL negotiation
- New feature: Support of RRPP for the link-up delay timer
- New feature: gRPC
- New feature: Protocol buffers
- New feature: Specifying the NTP time-offset thresholds for log and trap outputs
- New feature: Specifying the SNTP time-offset thresholds for log and trap outputs
- New feature: Ignoring the permit flag added by OpenFlow
- New feature: ND flood suppression
- New feature: IP address configuration of the VTEPs in a DR system
- New feature: Redistribution of MAC/IP advertisement routes into BGP unicast routing tables
- New feature: Route exchange between the BGP EVPN address family and the BGP VPNv4 or VPNv6 address family
- New feature: Disabling the ACLs issued by the OVSDB controller
- New feature: Replacement of the L3 VXLAN ID and RD of IP prefix advertisement routes on EDs
- New feature: BGP EVPN route advertisement suppression
- New feature: Removal of the default-gateway extended community attribute from EVPN gateway routes
- New feature: Configuring a 16/32-Gbps FC interface
- New feature: Enabling area ID-to-F_Port binding
- New feature: BB_Credit recovery for FC interfaces
- Modified feature: Software upgrade
- Modified feature: Using issu commands for an ISSU
- Modified feature: Per-packet load sharing mode of aggregation groups
- Modified feature: Configuring MAC-to-VLAN entries
- Modified feature: Configuring the advertisable TLVs
- Modified feature: Specifying the management address advertised in global management address TLV advertisement setting
- Modified feature: Configuring aging timer for dynamic ARP entries
- Modified feature: Enabling ARP snooping
- Modified feature: Displaying ARP snooping entries
- Modified feature: Clearing ARP snooping entries
- Modified feature: Configuring DHCP snooping handling strategy for Option 82 in request messages
- Modified feature: Setting the interface MTU for IPv6 packets
- Modified feature: Support for specifying next hop indexes in static routes
- Modified feature: Support for configuring OSPF commands in VSI interface view
- Modified feature: Displaying IS-IS LSP log information
- Modified feature: Clearing IS-IS LSP log information
- Modified feature: Filtering advertised/received BGP routes
- Modified feature: Filtering routes advertised to or received from a peer or peer group by using an ACL
- Modified feature: Configuring BMP

- Modified feature: Displaying detailed BGP routing information
- Modified feature: Applying a routing policy to routes outgoing to a peer or peer group
- Modified feature: Specifying an ACL
- Modified feature: Defining an ACL match criterion in a traffic class of a QoS policy
- Modified feature: Applying a QoS policy globally
- Modified feature: Configuring a test profile for RADIUS server status detection
- Modified feature: RADIUS server quiet timer
- Modified feature: Specifying the source IP address for outgoing RADIUS packets
- Modified feature: Specifying the source IP address for outgoing HWTACACS packets
- Modified feature: Including user IP addresses in MAC authentication requests
- Modified feature: Configuring MAC-based MAC authentication user accounts
- Modified feature: Handling new MAC access attempts in a VLAN after port security's MAC address limit for that VLAN is reached
- Modified feature: Port security NTK feature
- Modified feature: Password handling manners with password control enabled globally
- Modified feature: RSA key modulus length
- Modified feature: Key modulus length of the RSA key pair used for certificate request in a PKI domain
- Modified feature: Configuring the ECDSA signature authentication method in an IKE proposal
- Modified feature: Specifying the DH group used for key negotiation in IKE phase 1
- Modified feature: Keyboard-interactive authentication support for SSH users
- Modified feature: Displaying IPv4 source guard bindings
- Modified feature: Displaying IPv6 source guard bindings
- Modified feature: Configuring ARP attack detection logging
- Modified feature: Creating a BFD session for detecting the local interface state
- Modified feature: Configuring the BFD authentication mode for single-hop BFD control packets
- Modified feature: Setting the minimum interval for receiving BFD echo packets
- Modified feature: Support for specific BFD commands in VSI interface view
- Modified feature: Setting the maximum number of active routes in a VPN instance
- Modified feature: Enabling BFD on an MPLS TE tunnel interface
- Modified feature: Setting the length of ICMP or ICMPv6 echo requests for the ping operation
- Modified feature: Specifying a source interface for NTP messages
- Modified feature: PTP
- Modified feature: Shutting down an interface by using OpenFlow
- Modified feature: Configuring flow sampling of sFlow
- Modified feature: Configuring counter sampling of sFlow
- Modified feature: Testing the reachability of remote VMs in VXLANs
- Modified feature: Enabling packet statistics for automatically created VXLAN tunnels
- Modified feature: Disabling flooding for a VSI
- Modified feature: Setting the maximum bandwidth for an AC
- Modified feature: Displaying information about peers automatically discovered through BGP in EVPN
- Modified feature: PW redundancy

- [Modified feature: Displaying log buffer information and buffered logs](#)
- [Modified feature: Configuring the timestamp for logs sent to log hosts](#)
- [Modified feature: Configuring log output to a log host](#)
- [Modified feature: IRF master election during automated VCF fabric deployment](#)

New feature: Filtering the output from a display command by using multiple conditions

Filtering the output from a display command by using multiple conditions

You can execute the `display [| [by-linenum] { begin | exclude | include }]` command to use multiple conditions to filter the output from a `display` command.

Command reference

`display [| [by-linenum] { begin | exclude | include }]`

Use `display [| [by-linenum] { begin | exclude | include }]` to configure multiple conditions to filter the output from a `display` command.

Syntax

`display command [| [by-linenum] { begin | exclude | include } regular-expression]&<1-128>`

Views

Any view

Predefined user roles

network-admin
network-operator

Parameters

command: Specifies the keywords and arguments of a `display` command. To display available keywords and arguments, enter `display ?`.

`[| [by-linenum] { begin | exclude | include } regular-expression]&<1-128>`: Uses a maximum of 128 conditions to filter the output from the `display` command. Only the lines that meet all the conditions are displayed.

- **by-linenum**: Numbers the output lines. To number the output lines, you need to specify this keyword in only one of the filter conditions. If you do not specify this keyword in any of the filter conditions, the output lines are not numbered.
- **begin**: Displays the first line matching the specified regular expression and all subsequent lines.
- **exclude**: Displays all lines not matching the specified regular expression.
- **include**: Displays all lines matching the specified regular expression.
- *regular-expression*: Specifies a regular expression, a case-sensitive string of 1 to 256 characters.

Usage guidelines

The output from a **display** command might have many lines. You can use this command to filter and number the output to quickly locate the lines of interest. For more information about regular expressions, see CLI usage in *Fundamentals Configuration Guide*.

A line number is a 5-character string and might be followed by a colon (:) or hyphen (-). If you specify both **by-linenum** and **begin** *regular-expression* for a **display** command, a hyphen is displayed for all lines that do not match the regular expression.

Examples

Display log messages in the log buffer that contains both **SHELL** and **VTY**.

```
<Sysname> display logbuffer | include SEHLL | include VTY
%Sep 6 10:38:12:320 2018 Sysname SHELL/5/SHELL_LOGIN: VTY logged in from 169.254.100.171.
%Sep 6 10:52:32:576 2018 Sysname SHELL/5/SHELL_LOGOUT: VTY logged out from
169.254.100.171.
%Sep 6 16:03:27:100 2018 Sysname SHELL/5/SHELL_LOGIN: VTY logged in from 169.254.100.171.
%Sep 6 16:44:18:113 2018 Sysname SHELL/5/SHELL_LOGOUT: VTY logged out from
169.254.100.171.
```

New feature: Executing a batch file

Executing a batch file

About batch file and batch file execution

A batch file contains a set of commands. Executing a batch file executes the commands in the file one by one.

The system executes the commands in a batch file in the background and does not output the execution results. To verify the execution results, use the **display current-configuration** command.

Restrictions and guidelines

To execute a batch file on the device, create a batch file on a PC and load the batch file to the device.

As a best practice, try every command on the device to make sure the command line can be executed correctly before adding the command to a batch file. If a command is invalid or a condition for executing the command is not met, the command fails and the system continues to execute the next command.

Procedure

1. Enter system view.
system-view
2. Execute a batch file.
execute *filename*

Command reference

execute

Use **execute** to execute a batch file.

Syntax

execute *filename*

Views

System view

Predefined user roles

network-admin

Parameters

filename: Specifies the name of a batch file.

Usage guidelines

A batch file contains a set of commands. Executing a batch file executes the commands in the file one by one.

To execute a batch file on the device, create a batch file on a PC and load the batch file to the device.

As a best practice, try every command on the device to make sure the command line can be executed correctly before adding the command to a batch file. If a command is invalid or a condition for executing the command is not met, the command fails and the system continues to execute the next command.

The system executes the commands in a batch file in the background and does not output the execution results. To verify the execution results, use the **display current-configuration** command.

Examples

```
# Execute batch file test.bat.  
<Sysname> fdisk usba: 3  
[Sysname] execute test.bat
```

New feature: Disabling automatic system-wide next-startup configuration file operations

Disabling automatic system-wide next-startup configuration file operations

About automatic system-wide next-startup configuration file operations

By default, the system performs the following operations on all IRF subordinate devices in addition to the master device:

- Saves the running configuration to the next-startup configuration file on each member device when you execute the **save [safely] [backup | main] [force] [changed]** command.
- Deletes the next-startup configuration file on each member device when you execute the **reset saved-configuration** command.

When the system has a large amount of configuration, managing the configuration on all member devices will take a long time.

To have the system save the running configuration or delete the next-startup configuration file only on the master device, perform this task.

Disabling automatic system-wide next-startup configuration file operations will cause next-startup configuration inconsistency between the master and subordinate devices.

You can enable or disable the automatic system-wide next-startup configuration file operations, depending on your network requirements.

Procedure

1. Enter system view.
system-view
2. Disable automatic system-wide next-startup configuration file operations.
undo standby auto-update config

By default, next-startup configuration file operations are automatically synchronized across the entire system.

Command reference

standby auto-update config

Use **standby auto-update config** to enable automatic system-wide next-startup configuration file operations.

Use **undo standby auto-update config** to disable automatic system-wide next-startup configuration file operations.

Syntax

standby auto-update config

undo standby auto-update config

Default

Next-startup configuration file operations are automatically synchronized across the entire system.

Views

System view

Predefined user roles

network-admin

Usage guidelines

By default, the system performs the following operations on all IRF subordinate devices in addition to the master device:

- Saves the running configuration to the next-startup configuration file on each member device when you execute the **save [safely] [backup | main] [force] [changed]** command.
- Deletes the next-startup configuration file on each member device when you execute the **reset saved-configuration** command.

When the system has a large amount of configuration, managing the configuration on all member devices will take a long time.

To have the system save the running configuration or delete the next-startup configuration file only on the master device, use the **undo standby auto-update config** command.

Disabling automatic system-wide next-startup configuration file operations will cause next-startup configuration inconsistency between the master and subordinate devices.

You can enable or disable the automatic system-wide next-startup configuration file operations, depending on your network requirements.

Examples

Enable automatic system-wide next-startup configuration file operations.

```
<Sysname> system-view
[Sysname] standby auto-update config
```

Related commands

```
reset saved-configuration
save
```

New feature: Using an encrypted configuration file to roll back configuration

Using an encrypted configuration file to roll back configuration

You can specify an encrypted configuration file for the **configuration replace file** command to roll back the configuration.

Command reference

Modified command: configuration replace file

Syntax

```
configuration replace file filename
```

Views

System view

Change description

Before modification: You cannot specify an encrypted configuration file for the *filename* argument.

After modification: You can specify an encrypted configuration file for the *filename* argument.

New feature: Support for encrypted configuration files for configuration comparison

Comparing encrypted configuration files with other configurations for their differences

You can compare encrypted configuration files with other configurations for their differences by using the **display current-configuration diff** and **display diff** commands.

Command reference

Modified command: display current-configuration diff

Syntax

```
display current-configuration diff
```

Views

Any view

Change description

Before modification: This command cannot compare an encrypted startup configuration file with the running configuration.

After modification: This command can compare an encrypted startup configuration file with the running configuration.

Modified command: display diff

Syntax

```
display diff configfile file-name-s { configfile file-name-d |  
current-configuration | startup-configuration }  
  
display diff current-configuration { configfile file-name-d |  
startup-configuration }  
  
display diff startup-configuration { configfile file-name-d |  
current-configuration }
```

Views

Any view

Change description

Before modification: This command cannot compare an encrypted configuration file with other configurations for their differences.

After modification: This command can compare an encrypted configuration file with other configurations for their differences. You can use the **startup-configuration** keyword to specify the startup configuration file for comparison even if that file has been encrypted. In addition, you can use the *file-name-s* or *file-name-d* argument to specify an encrypted configuration file for comparison.

New feature: ND MAD on management Ethernet ports

Configuring ND MAD on management Ethernet ports

About ND MAD on management Ethernet ports

ND MAD detects multi-active collisions by using NS packets to transmit the IRF domain ID and the active ID (the member ID of the master).

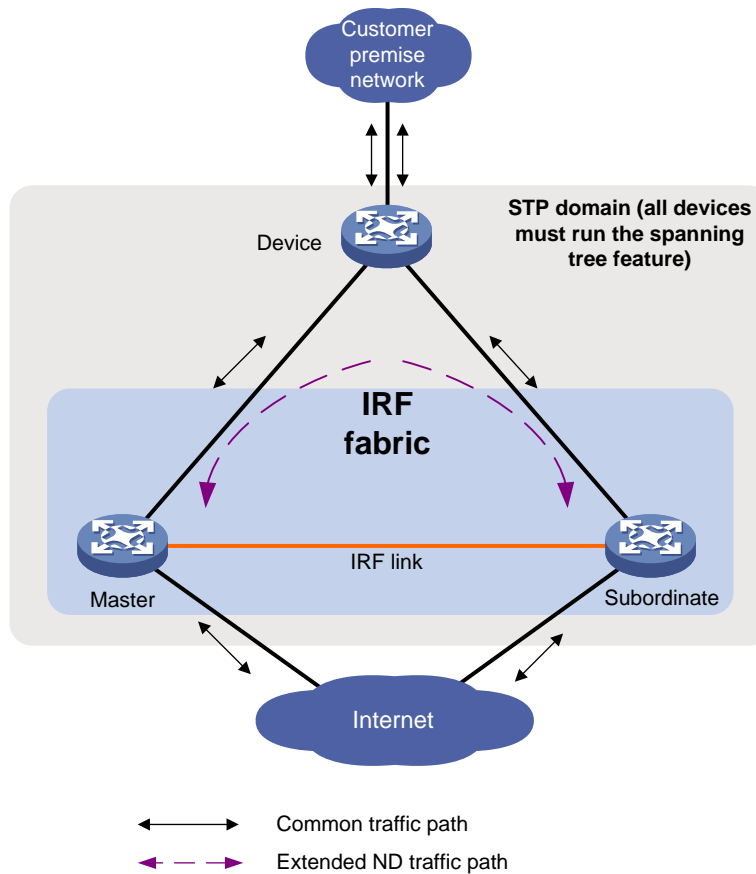
If management Ethernet ports are used, ND MAD must work with an intermediate device. Make sure the following requirements are met:

- Connect a management Ethernet port on each member device to the intermediate device.
- On the intermediate device, you must assign the ports used for ND MAD to the same VLAN.

Each IRF member device compares the domain ID and the active ID in incoming NS packets with its domain ID and active ID.

- If the domain IDs are different, the NS packet is from a different IRF fabric. The device does not continue to process the packet with the MAD mechanism.
- If the domain IDs are the same, the device compares the active IDs.
 - If the active IDs are different, the IRF fabric has split.
 - If the active IDs are the same, the IRF fabric is integrated.

Figure 1 ND MAD scenario



Restrictions and guidelines

When you configure ND MAD that uses management Ethernet ports, follow these restrictions and guidelines:

| Category | Restrictions and guidelines |
|--------------------------------------|---|
| Management Ethernet ports for ND MAD | Connect a management Ethernet port on each member device to the common Ethernet ports on the intermediate device. |
| ND MAD VLAN | On the intermediate device, create a VLAN for ND MAD, and assign the ports used for ND MAD to the VLAN. On the IRF fabric, you do not need to assign the management Ethernet ports to the VLAN. |
| ND MAD and feature configuration | <ul style="list-style-type: none"> Enable the IRF fabric to change its bridge MAC address as soon as the address owner leaves. If the intermediate device is also an IRF fabric, assign the two IRF fabrics different domain IDs for correct split detection. |

Procedure

- Enter system view.
system-view
- Assign a domain ID to the IRF fabric.
irf domain domain-id
The default IRF domain ID is 0.
- Configure the IRF bridge MAC address to change as soon as the address owner leaves.

undo irf mac-address persistent

By default, the IRF bridge MAC address remains unchanged for 12 minutes after the address owner leaves.

4. Enter management Ethernet interface view.

interface m-gigabitethernet *interface-number*

Of all management Ethernet ports on an IRF fabric, only the master's management Ethernet port is accessible.

5. Assign an IPv6 address to the management Ethernet port.

ipv6 address { *ipv6-address/pre-length* | *ipv6 address pre-length* }

By default, no IPv6 addresses are assigned to a management Ethernet port.

6. Enable ND MAD.

mad nd enable

By default, ND MAD is disabled.

Command reference

mad nd enable

Use **mad nd enable** to enable ND MAD.

Use **undo mad nd enable** to disable ND MAD.

Syntax

mad nd enable

undo mad nd enable

Default

ND MAD is disabled.

Views

Management Ethernet interface view

Predefined user roles

network-admin

Usage guidelines

Do not configure ND MAD together with LACP MAD, BFD MAD, or ARP MAD, because they handle collisions differently.

When you configure ND MAD that uses management Ethernet ports, follow these restrictions and guidelines:

| Category | Restrictions and guidelines |
|--------------------------------------|--|
| Management Ethernet ports for ND MAD | Connect a management Ethernet port on each member device to the common Ethernet ports on the intermediate device. |
| ND MAD VLAN | On the intermediate device, create a VLAN for ND MAD, and assign the ports used for ND MAD to the VLAN. On the IRF fabric, you do not need to assign the management Ethernet ports to the VLAN. |
| ND MAD and feature configuration | <ul style="list-style-type: none">• Enable the IRF fabric to change its bridge MAC address as soon as the address owner leaves.• If the intermediate device is also an IRF fabric, assign the two IRF fabrics different domain IDs for correct split detection. |

When you use the **mad nd enable** command, the system prompts you to enter a domain ID. If you do not want to change the current domain ID, press **enter** at the prompt.

An IRF fabric has only one IRF domain ID. You can change the IRF domain ID by using the following commands: **irf domain**, **mad enable**, **mad arp enable**, or **mad nd enable**. The IRF domain IDs configured by using these commands overwrite each other.

Examples

```
# Enable ND MAD on M-GigabitEthernet 0/0/1.
<Sysname> system-view
[Sysname] interface m-gigabitethernet 0/0/1
[Sysname-M-GigabitEthernet0/0/1] mad nd enable
You need to assign a domain ID (range: 0-4294967295)
[Current domain is: 0]: 1
The assigned domain ID is: 1
```

Related commands

irf domain

New feature: Fast retrain

Configuring fast retrain

About fast retrain

Copper Ethernet ports negotiate parameters through electric signals to determine the up/down status. In an actual network, if a copper port is in a high-frequency or low-frequency environment, the electrical signals will be affected. As a result, the negotiated status might be incorrect. You can enable or disable fast retrain on a copper port to adjust the frequency of electrical signals transmitted by the copper port to improve the anti-attenuation and anti-interference capabilities of the interface.

Enabling fast retrain on a copper port enables the port to perform negotiation by using high-frequency electrical signals, which improves the anti-attenuation capability of the copper port. Disabling fast retrain on a copper port enables the port to perform negotiation by using low-frequency electrical signals, which improves the anti-interference capability of the copper port.

Restrictions and guidelines

This feature is supported only on copper ports of the LSWM124XGT2Q (JH182A) interface card of the following switches:

- HPE FlexFabric 5940 2-slot Switch (JH397A).
- HPE FlexFabric 5940 4-slot Switch (JH398A).

Procedure

1. Enter system view.
system-view
2. Enter Ethernet interface view.
interface *interface-type interface-number*
3. Enable fast retrain.
port fast-retrain enable
By default, fast retrain is disabled.

Command reference

port fast-retrain enable

Use **port fast-retrain enable** to enable fast retrain.

Use **undo port fast-retrain enable** to disable fast retrain.

Syntax

```
port fast-retrain enable
undo port fast-retrain enable
```

Default

Fast retrain is disabled.

Views

Ethernet interface view

Predefined user roles

network-admin

Usage guidelines

This command is supported only on copper ports of the LSWM124XGT2Q (JH182A) interface card of the following switches:

- HPE FlexFabric 5940 2-slot Switch (JH397A).
- HPE FlexFabric 5940 4-slot Switch (JH398A).

Copper Ethernet ports negotiate parameters through electric signals to determine the up/down status. In an actual network, if a copper port is in a high-frequency or low-frequency environment, the electrical signals will be affected. As a result, the negotiated status might be incorrect. You can enable or disable fast retrain on a copper port to adjust the frequency of electrical signals transmitted by the copper port to improve the anti-attenuation and anti-interference capabilities of the interface.

Enabling fast retrain on a copper port enables the port to perform negotiation by using high-frequency electrical signals, which improves the anti-attenuation capability of the copper port. Disabling fast retrain on a copper port enables the port to perform negotiation by using low-frequency electrical signals, which improves the anti-interference capability of the copper port.

Examples

```
# Disable fast retrain on Ten-GigabitEthernet 1/0/1.
```

```
<Sysname> system-view
```

```
[Sysname] interface ten-gigabitethernet 1/0/1
```

```
[Sysname-Ten-GigabitEthernet1/0/1] undo port fast-retrain enable
```

New feature: Displaying the status and packet statistics of interfaces

Displaying the status and packet statistics of interfaces

This feature allows you to display the status and packet statistics of interfaces.

Command reference

display interface link-info

Use **display interface link-info** to display the status and packet statistics of interfaces.

Syntax

```
display interface link-info [ main ]
```

Views

Any view

Predefined user roles

network-admin

Parameters

main: Specifies all interfaces except subinterfaces. If you do not specify this keyword, this command displays status and packet statistics of all interfaces.

Examples

Display status and statistics of all interfaces.

```
<Sysname> display interface link-info
```

Link: ADM - administratively down; Stby - standby

Protocol: (s) - spoofing

| Interface | Link | Protocol | InUsage | OutUsage | InErrs | OutErrs |
|-----------|------|----------|---------|----------|--------|---------|
| XGE1/0/1 | UP | UP | 10.09% | 0% | 0 | 0 |
| NULL0 | UP | UP(s) | 0% | 0% | 0 | 0 |

Overflow: More than 7 digits.

--: Not supported.

Table 19 Command output

| Field | Description |
|---|---|
| Link: ADM - administratively down; Stby - standby | Physical link state of the interface: <ul style="list-style-type: none">ADM—The interface has been shut down by using the shutdown command. To restore the physical state of the interface, use the undo shutdown command.Stby—The interface is a backup interface in standby state. To see the primary interface, use the display interface-backup state command. |
| Protocol: (s) – spoofing | The data link layer protocol of the interface is up, but the link is an on-demand link or does not exist. The (s) attribute represents the spoofing flag. This value is typical of null interfaces, loopback interfaces, and InLoopback interfaces. |
| Interface | Abbreviated interface name. |
| Link | Physical link state of the interface: <ul style="list-style-type: none">UP—The interface is physically up.DOWN—The interface is physically down.ADM—The interface has been shut down by using the shutdown command. To restore the physical state of the interface, use the undo shutdown command.Stby—The interface is a backup interface in standby state. |

| Field | Description |
|-------------------------------|--|
| Protocol | Data link layer protocol state of the interface: <ul style="list-style-type: none"> UP—The data link layer protocol of the interface is up. DOWN—The data link layer protocol of the interface is down. UP(s)—The data link layer protocol of the interface is up, but the link is an on-demand link or does not exist. The (s) attribute represents the spoofing flag. This value is typical of null interfaces, loopback interfaces, and InLoopback interfaces. |
| InUsage | Inbound bandwidth usage within the most recent statistics polling interval. It is calculated by this formula: Average inbound speed of the interface within the most recent statistics polling interval/interface bandwidth. When the usage is smaller than 0.01%, 0.01% is displayed. To set the statistics polling interval, use the flow-interval command. |
| OutUsage | Outbound bandwidth usage within the most recent statistics polling interval. It is calculated by this formula: Average outbound speed of the interface within the most recent statistics polling interval/interface bandwidth. When the usage is smaller than 0.01%, 0.01% is displayed. To set the statistics polling interval, use the flow-interval command. |
| InErrs | Number of error packets received. |
| OutErrs | Number of error packets sent. |
| Overflow: More than 7 digits. | The data length of a statistical item value is greater than 7 decimal digits. |
| --: Not supported. | A hyphen (-) indicates that the corresponding statistical item is not supported. |

Related commands

flow-interval

New feature: Displaying operating status and information of all interfaces except subinterfaces

Displaying operating status and information of all interfaces except subinterfaces

This feature allows you to display operating status and information of all interfaces except subinterfaces.

Command reference

display interface main

Use **display interface main** to display operating status and information of all interfaces except subinterfaces.

Syntax

display interface [*interface-type*] [**brief** [**description** | **down**]] **main**

Views

Any view

Predefined user roles

network-admin

Parameters

interface-type: Specifies an interface type. If you do not specify this argument, the command displays information about interfaces of all types.

brief: Displays brief interface information. If you do not specify this keyword, the command displays detailed interface information.

description: Displays complete interface descriptions. If you do not specify this keyword, the command displays only the first 27 characters of each interface description.

down: Displays information about interfaces in down state and the causes. If you do not specify this keyword, the command displays information about interfaces in all states.

Examples

Display operating status and information of all interfaces except subinterfaces.

```
<Sysname> display interface main
Ten-GigabitEthernet1/0/1
Current state: Administratively DOWN
Line protocol state: DOWN
IP packet frame type: Ethernet II, hardware address: 4005-6538-0100
Description: Ten-GigabitEthernet1/0/1 Interface
Bandwidth: 1000000 kbps
Loopback is not set
Unknown-speed mode, unknown-duplex mode
Link speed type is autonegotiation, link duplex type is autonegotiation
Flow-control is not enabled
Maximum frame length: 9216
Allow jumbo frames to pass
Broadcast max-ratio: 100%
Multicast max-ratio: 100%
Unicast max-ratio: 100%
PVID: 1
MDI type: Automdix
Port link-type: Access
    Tagged VLANs:    None
    Untagged VLANs: 1
Port priority: 2
Last link flapping: Never
Last clearing of counters: Never
Current system time:2018-04-11 10:20:24
Last time when physical state changed to up:-
Last time when physical state changed to down:2018-04-11 09:11:09
    Peak input rate: 0 bytes/sec, at 00-00-00 00:00:00
    Peak output rate: 0 bytes/sec, at 00-00-00 00:00:00
    Last 300 second input: 0 packets/sec 0 bytes/sec -%
    Last 300 second output: 0 packets/sec 0 bytes/sec -%
    Input (total):  0 packets, 0 bytes
                    0 unicasts, 0 broadcasts, 0 multicasts, 0 pauses
    Input (normal): 0 packets, 0 bytes
```

```

        0 unicasts, 0 broadcasts, 0 multicasts, 0 pauses
Input:  0 input errors, 0 runts, 0 giants, 0 throttles
        0 CRC, 0 frame, 0 overruns, 0 aborts
        0 ignored, 0 parity errors
Output (total): 0 packets, 0 bytes
        0 unicasts, 0 broadcasts, 0 multicasts, 0 pauses
Output (normal): 0 packets, 0 bytes
        0 unicasts, 0 broadcasts, 0 multicasts, 0 pauses
Output:  0 output errors, 0 underruns, 0 buffer failures
        0 aborts, 0 deferred, 0 collisions, 0 late collisions
        0 lost carrier, 0 no carrier

Ten-GigabitEthernet1/0/2
Current state: Administratively DOWN
Line protocol state: DOWN
Description: Ten-GigabitEthernet1/0/2 Interface
Bandwidth: 1000000 kbps
Flow-control is not enabled
Maximum transmission unit: 1500
Allow jumbo frames to pass
Broadcast max-ratio: 100%
Multicast max-ratio: 100%
Unicast max-ratio: 100%
Internet protocol processing: Disabled
IP packet frame type: Ethernet II, hardware address: 4005-6538-0107
IPv6 packet frame type: Ethernet II, hardware address: 4005-6538-0107
Output queue - Urgent queuing: Size/Length/Discards 0/100/0
Output queue - Protocol queuing: Size/Length/Discards 0/500/0
Output queue - FIFO queuing: Size/Length/Discards 0/75/0
Last link flapping: Never
Last clearing of counters: Never
Current system time:2018-04-11 10:20:24
Last time when physical state changed to up:-
Last time when physical state changed to down:2018-04-11 09:33:52
Peak input rate: 0 bytes/sec, at 00-00-00 00:00:00
Peak output rate: 0 bytes/sec, at 00-00-00 00:00:00
Last 300 second input: 0 packets/sec 0 bytes/sec -%
Last 300 second output: 0 packets/sec 0 bytes/sec -%
Input (total):  0 packets, 0 bytes
        0 unicasts, 0 broadcasts, 0 multicasts, 0 pauses
Input (normal):  0 packets, 0 bytes
        0 unicasts, 0 broadcasts, 0 multicasts, 0 pauses
Input:  0 input errors, 0 runts, 0 giants, 0 throttles
        0 CRC, 0 frame, 0 overruns, 0 aborts
        0 ignored, 0 parity errors
Output (total): 0 packets, 0 bytes
        0 unicasts, 0 broadcasts, 0 multicasts, 0 pauses
Output (normal): 0 packets, 0 bytes

```

0 unicasts, 0 broadcasts, 0 multicasts, 0 pauses
Output: 0 output errors, 0 underruns, 0 buffer failures
0 aborts, 0 deferred, 0 collisions, 0 late collisions
0 lost carrier, 0 no carrier

Table 20 Command output

| Field | Description |
|---------------------|---|
| Current state | <p>Physical link state of the interface:</p> <ul style="list-style-type: none"> • Administratively DOWN—The interface has been shut down by using the shutdown command. • DOWN—The interface is administratively up, but its physical state is down (possibly because no physical link exists or the link has failed). • DOWN (Link-Aggregation interface down)—The aggregate interface to which the interface belongs has been shut down by using the shutdown command. • DOWN (Tunnel-Bundle administratively down)—The tunnel bundle interface to which the interface belongs has been shut down by using the shutdown command. • Link-Flap DOWN—The interface has been shut down by the link flapping protection feature. • ETH-rddc Shutdown—The interface has been shut down by the Reth module. • IRF-link-down—The interface has been shut down by IRF. This state occurs when the IRF member device that contains the interface is detected to have lost the connectivity of all its IRF links in the MDC. • mac-address moving down—The interface has been shut down by the MAC address move suppression feature. • MAD ShutDown—The interface has been shut down by IRF MAD. This state occurs if the interface is on an IRF fabric placed in Recovery state after an IRF split. • OFF DOWN—The interface has been shut down by OpenFlow. • Storm-Constrain—The interface has been shut down because the storm control feature detected that unknown unicast traffic, multicast traffic, or broadcast traffic exceeded the upper threshold. • STP DOWN—The interface has been shut down by the BPDU guard feature. • UP—The interface is both administratively and physically up. |
| Line protocol state | <p>Data link layer state of the interface. The state is determined through automatic parameter negotiation at the data link layer.</p> <ul style="list-style-type: none"> • UP—The data link layer protocol is up. • UP (spoofing)—The data link layer protocol is up, but the link is an on-demand link or does not exist. This attribute is typical of null interfaces and loopback interfaces. • DOWN—The data link layer protocol is down. • DOWN (protocols)—The data link layer has been shut down by protocols included in the parentheses. Available protocols include: <ul style="list-style-type: none"> ◦ DLDP—Shuts down the data link layer when it detects that the link is unidirectional. |

| Field | Description |
|--|--|
| | <ul style="list-style-type: none"> ○ OAM—Shuts down the data link layer when it detects a remote link failure. ○ LAGG—Shuts down the data link layer when it detects that the aggregate interface does not have Selected ports. ○ BFD—Shuts down the data link layer when it detects a link failure. ○ MACSEC—Shuts down the data link layer when it fails to negotiate the encryption parameters. |
| Bandwidth | Expected bandwidth of the interface. |
| Maximum transmission unit | MTU of the interface. |
| Internet protocol processing: Disabled | The interface is not assigned an IP address and cannot process IP packets. |
| Internet address | IP address of the interface. The primary attribute indicates that the address is the primary IP address. |
| IP packet frame type | IPv4 packet framing format. |
| hardware address | MAC address of the interface. |
| IPv6 packet frame type | IPv6 packet framing format. |
| Port priority | Port priority of the interface. |
| Loopback is set internal | An internal loopback test is running on the interface. This field depends on your configuration. |
| Loopback is set external | An external loopback test is running on the interface. This field depends on your configuration. |
| Loopback is not set | No loopback test is running on the interface. This field depends on your configuration. |
| 10Mbps-speed mode | The interface is operating at 10 Mbps. This field depends on your configuration and the link parameter negotiation result. |
| 100Mbps-speed mode | The interface is operating at 100 Mbps. This field depends on your configuration and the link parameter negotiation result. |
| 1000Mbps-speed mode | The interface is operating at 1000 Mbps. This field depends on your configuration and the link parameter negotiation result. |
| 10Gbps-speed mode | The interface is operating at 10 Gbps. This field depends on your configuration and the link parameter negotiation result. |
| 40Gbps-speed mode | The interface is operating at 40 Gbps. This field depends on your configuration and the link parameter negotiation result. |
| 100Gbps-speed mode | The interface is operating at 100 Gbps. This field depends on your configuration and the link parameter negotiation result. |
| Unknown-speed mode | The speed of the interface is unknown because the speed negotiation fails or the interface is physically disconnected. |
| half-duplex mode | The interface is operating in half duplex mode. This field depends on your configuration and the link parameter negotiation result. |
| full-duplex mode | The interface is operating in full duplex mode. This field depends on your configuration and the link parameter negotiation result. |
| unknown-duplex mode | The duplex mode of the interface is unknown because the duplex mode negotiation fails or the interface is physically disconnected. |
| Link speed type is autonegotiation | The interface is configured with the speed auto command. |

| Field | Description |
|---------------------------------------|---|
| Link speed type is force link | The interface is manually configured with a speed (for example, 1000 Mbps) by using the speed command. |
| link duplex type is autonegotiation | The interface is configured with the duplex auto command. |
| link duplex type is force link | The interface is manually configured with a duplex mode (for example, half or full) by using the duplex command. |
| Flow-control is not enabled | Generic flow control is disabled on the interface. This field depends on your configuration and the link parameter negotiation result. |
| Maximum frame length | Maximum length of Ethernet frames allowed to pass through the interface. |
| Allow jumbo frame to pass | The interface allows jumbo frames to pass through. |
| Broadcast max- | Broadcast storm suppression threshold in ratio, pps, or kbps. The unit of the threshold depends on your configuration. |
| Multicast max- | Multicast storm suppression threshold in ratio, pps, or kbps. The unit of the threshold depends on your configuration. |
| Unicast max- | Unknown unicast storm suppression threshold in ratio, pps, or kbps. The unit of the threshold depends on your configuration. |
| PVID | Port VLAN ID (PVID) of the interface. |
| MDI type | MDIX mode of the interface: <ul style="list-style-type: none"> • automdix. • mdi. • mdix. |
| Port link-type | Link type of the interface: <ul style="list-style-type: none"> • access. • trunk. • hybrid. |
| Tagged VLANs | VLANs for which the interface sends packets without removing VLAN tags. |
| Untagged VLANs | VLANs for which the interface sends packets after removing VLAN tags. |
| VLAN Passing | VLANs whose packets can be forwarded by the port. The VLANs must have been created. |
| VLAN permitted | VLANs whose packets are permitted by the port. |
| Trunk port encapsulation | Encapsulation protocol type for the trunk port. |
| Last link flapping | The amount of time that has elapsed since the most recent physical state change of the interface. This field displays Never if the interface has been physically down since device startup. |
| Last clearing of counters | Time when the reset counters interface command was last used to clear the interface statistics. This field displays Never if the reset counters interface command has never been used on the interface since device startup. |
| Current system time | Current system time in the YYYY/MM/DD HH:MM:SS format. If the time zone is configured, this field is in the YYYY/MM/DD HH:MM:SS <i>zone-name</i> ±HH:MM:SS format, where the <i>zone-name</i> argument is the local time zone. |
| Last time when physical state changed | Last time when physical state of the interface changed to up. |

| Field | Description |
|---|---|
| to up | <p>If the time zone is configured, this field is in the YYYY/MM/DD HH:MM:SS <i>zone-name</i>±HH:MM:SS format, where the <i>zone-name</i> argument is the local time zone.</p> <p>A hyphen (-) indicates that the physical state of the interface has not changed to up.</p> |
| Last time when physical state changed to down | <p>Last time when physical state of the interface changed to down.</p> <p>If the time zone is configured, this field is in the YYYY/MM/DD HH:MM:SS <i>zone-name</i>±HH:MM:SS format, where the <i>zone-name</i> argument is the local time zone.</p> <p>A hyphen (-) indicates that the physical state of the interface has not changed to down.</p> |
| Last 300 second input: 0 packets/sec 0 bytes/sec 0% Last 300 second output: 0 packets/sec 0 bytes/sec 0% | <p>Average inbound or outbound traffic rate (in pps and Bps) in the last 300 seconds, and the ratio of the actual rate to the interface bandwidth.</p> <p>A hyphen (-) indicates that the statistical item is not supported.</p> |
| Input(total): 0 packets, 0 bytes 0 unicasts, 0 broadcasts, 0 multicasts, 0 pauses | <p>The two fields on the first line represent the inbound traffic statistics (in packets and bytes) for the interface. All inbound normal packets, abnormal packets, and normal pause frames were counted.</p> <p>The four fields on the second line represent:</p> <ul style="list-style-type: none"> • Number of inbound unicast packets. • Number of inbound broadcasts. • Number of inbound multicasts. • Number of inbound pause frames. <p>A hyphen (-) indicates that the statistical item is not supported.</p> |
| Input(normal): 0 packets, 0 bytes 0 unicasts, 0 broadcasts, 0 multicasts, 0 pauses | <p>The two fields on the first line represent the inbound normal traffic and pause frame statistics (in packets and bytes) for the interface.</p> <p>The four fields on the second line represent:</p> <ul style="list-style-type: none"> • Number of inbound normal unicast packets. • Number of inbound normal broadcasts. • Number of inbound normal multicasts. • Number of inbound normal pause frames. <p>A hyphen (-) indicates that the statistical item is not supported.</p> |
| input errors | Statistics of incoming error packets. |
| runts | <p>Number of inbound frames meeting the following conditions:</p> <ul style="list-style-type: none"> • Shorter than 64 bytes. • In correct format. • Containing valid CRCs. |
| giants | <p>Number of inbound giants. Giants refer to frames larger than the maximum frame length supported on the interface.</p> <p>For an Ethernet interface that does not permit jumbo frames, the maximum frame length is as follows:</p> <ul style="list-style-type: none"> • 1518 bytes (without VLAN tags). • 1522 bytes (with VLAN tags). <p>For an Ethernet interface that permits jumbo frames, the maximum Ethernet frame length is set when you configure jumbo frame support on the interface.</p> |
| throttles | Number of inbound frames that had a non-integer number of bytes. |

| Field | Description |
|--|---|
| CRC | Total number of inbound frames that had a normal length, but contained CRC errors. |
| frame | Total number of inbound frames that contained CRC errors and a non-integer number of bytes. |
| overruns | Number of packets dropped because the input rate of the port exceeded the queuing capability. |
| aborts | <p>Total number of illegal inbound packets:</p> <ul style="list-style-type: none"> • Fragment frames—CRC error frames shorter than 64 bytes. The length (in bytes) can be an integral or non-integral value. • Jabber frames—CRC error frames greater than the maximum frame length supported on the Ethernet interface (with an integral or non-integral length). <ul style="list-style-type: none"> ○ For an Ethernet interface that does not permit jumbo frames, the maximum frame length is 1518 bytes (without VLAN tags) or 1522 bytes (with VLAN tags). ○ For an Ethernet interface that permits jumbo frames, the maximum Ethernet frame length is set when you configure jumbo frame support on the interface. • Symbol error frames—Frames that contained a minimum of one undefined symbol. • Unknown operation code frames—Non-pause MAC control frames. • Length error frames—Frames whose 802.3 length fields did not match the actual frame length (46 to 1500 bytes). |
| ignored | Number of inbound frames dropped because the receiving buffer of the port ran low. |
| parity errors | Total number of frames with parity errors. |
| Output(total): 0 packets, 0 bytes 0 unicasts, 0 broadcasts, 0 multicasts, 0 pauses | <p>The two fields on the first line represent the outbound traffic statistics (in packets and bytes) for the interface. All outbound normal packets, abnormal packets, and normal pause frames were counted.</p> <p>The four fields on the second line represent:</p> <ul style="list-style-type: none"> • Number of outbound unicast packets. • Number of outbound broadcasts. • Number of outbound multicasts. • Number of outbound pause frames. <p>A hyphen (-) indicates that the statistical item is not supported.</p> |
| Output(normal): 0 packets, 0 bytes 0 unicasts, 0 broadcasts, 0 multicasts, 0 pauses | <p>The two fields on the first line represent the outbound normal traffic and pause frame statistics (in packets and bytes) for the interface.</p> <p>The four fields on the second line represent:</p> <ul style="list-style-type: none"> • Number of outbound normal unicast packets. • Number of outbound normal broadcasts. • Number of outbound normal multicasts. • Number of outbound normal pause frames. <p>A hyphen (-) indicates that the statistical item is not supported.</p> |
| output errors | Number of outbound packets with errors. |
| underruns | Number of packets dropped because the output rate of the interface exceeded the output queuing capability. This is a |

| Field | Description |
|------------------|--|
| | low-probability hardware anomaly. |
| buffer failures | Number of packets dropped because the transmitting buffer of the interface ran low. |
| aborts | Number of packets that failed to be transmitted, for example, because of Ethernet collisions. |
| deferred | Number of frames that the interface deferred to transmit because of detected collisions. |
| collisions | Number of frames that the interface stopped transmitting because Ethernet collisions were detected during transmission. |
| late collisions | Number of frames that the interface deferred to transmit after transmitting their first 512 bits because of detected collisions. |
| lost carrier | Number of carrier losses during transmission. This counter increases by one when a carrier is lost, and applies to serial WAN interfaces. |
| no carrier | Number of times that the port failed to detect the carrier when attempting to send frames. This counter increases by one when a port failed to detect the carrier, and applies to serial WAN interfaces. |
| Peak input rate | Peak rate of inbound traffic in Bps, and the time when the peak inbound traffic rate occurred. |
| Peak output rate | Peak rate of outbound traffic in Bps, and the time when the peak outbound traffic rate occurred. |

Display brief information of all interfaces except subinterfaces.

```
<Sysname> display interface brief main
Brief information on interfaces in route mode:
Link: ADM - administratively down; Stby - standby
Protocol: (s) - spoofing
```

| Interface | Link | Protocol | Primary IP | Description |
|-----------|------|----------|--------------|-------------|
| XGE1/0/1 | DOWN | DOWN | -- | |
| Loop0 | UP | UP(s) | 2.2.2.9 | |
| NULL0 | UP | UP(s) | -- | |
| Vlan1 | UP | DOWN | -- | |
| Vlan999 | UP | UP | 192.168.1.42 | |

```
Brief information on interfaces in bridge mode:
Link: ADM - administratively down; Stby - standby
Speed: (a) - auto
```

```
Duplex: (a)/A - auto; H - half; F - full
Type: A - access; T - trunk; H - hybrid
```

| Interface | Link | Speed | Duplex | Type | PVID | Description |
|-----------|------|-------|--------|------|------|-------------|
| XGE1/0/2 | DOWN | auto | A | A | 1 | |

Display brief information about all interfaces, including the complete interface descriptions.

```
<Sysname> display interface brief description main
Brief information on interfaces in bridge mode:
Link: ADM - administratively down; Stby - standby
```

```

Speed: (a) - auto
Duplex: (a)/A - auto; H - half; F - full
Type: A - access; T - trunk; H - hybrid
Interface          Link Speed  Duplex Type PVID Description
XGE1/0/3           UP   auto   F(a)  A    1    aaaaaaaaaaaaaaaaaaaaaaaaaaaaaa
aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa

```

Display information about interfaces in DOWN state and the causes.

```

<Sysname> display interface brief down main
Brief information on interfaces in route mode:
Link: ADM - administratively down; Stby - standby
Interface          Link Cause
GE1/0/1            DOWN Not connected
Vlan2              DOWN Not connected

```

```

Brief information on interfaces in bridge mode:
Link: ADM - administratively down; Stby - standby
Interface          Link Cause
XGE1/0/2           DOWN Not connected

```

Table 21 Command output

| Field | Description |
|---|---|
| Brief information on interfaces in route mode: | Brief information about Layer 3 interfaces. |
| Interface | Interface name. |
| Link | Physical link state of the interface: <ul style="list-style-type: none"> UP—The interface is physically up. DOWN—The interface is physically down. ADM—The interface has been shut down by using the shutdown command. To restore the physical state of the interface, use the undo shutdown command. Stby—The interface is a backup interface in standby state. To see the primary interface, use the display interface-backup state command. |
| Protocol | Data link layer protocol state of the interface: <ul style="list-style-type: none"> UP—The data link layer protocol of the interface is up. DOWN—The data link layer protocol of the interface is down. UP(s)—The data link layer protocol of the interface is up, but the link is an on-demand link or does not exist. The (s) attribute represents the spoofing flag. This value is typical of null interfaces and loopback interfaces. |
| Primary IP | Primary IP address of the interface. This field displays two hyphens (--) if the interface does not have an IP address. |
| Description | Description of the interface. |
| Brief information of interfaces in bridge mode: | Brief information about Layer 2 interfaces. |
| Type: A - access; T - trunk; H - hybrid | Link type options for interfaces. |
| Speed | Speed of the interface, in bps. |

| Field | Description |
|--------|--|
| | <p>This field displays the (a) flag next to the speed if the speed is automatically negotiated.</p> <p>This field displays auto if the interface is configured to autonegotiate its speed but the autonegotiation has not started.</p> |
| Duplex | <p>Duplex mode of the interface:</p> <ul style="list-style-type: none"> • A—Autonegotiation. The interface is configured to autonegotiate its duplex mode but the autonegotiation has not started. • F—Full duplex. • F(a)—Autonegotiated full duplex. • H—Half duplex. • H(a)—Autonegotiated half duplex. |
| Type | <p>Link type of the interface:</p> <ul style="list-style-type: none"> • A—Access. • H—Hybrid. • T—Trunk. |
| PVID | Port VLAN ID. |
| Cause | <p>Cause for the physical link state of an interface to be DOWN:</p> <ul style="list-style-type: none"> • Administratively—The interface has been manually shut down by using the shutdown command. To restore the physical state of the interface, use the undo shutdown command. • DOWN (Link-Aggregation interface down)—The interface is a member port of an aggregate interface, and the aggregate interface is down. • DOWN (Loopback detection down)—The loopback detection module has detected loops. • DOWN (Monitor-Link uplink down)—The monitor link module has detected that the uplink is down. • IRF-link-down—The IRF member device that contains the interface has lost the connectivity of all its IRF links in the MDC. • MAD ShutDown—The interface is on an IRF fabric placed by IRF MAD in Recovery state after an IRF split. • Not connected—No physical connection exists (possibly because the network cable is disconnected or faulty). • Storm-Constrain—The storm control feature has detected that unknown unicast traffic, multicast traffic, or broadcast traffic exceeded the upper threshold. • STP DOWN—The interface has been shut down by the BPDU guard feature. • Port Security Disabled—The interface has been shut down by the intrusion detection mechanism because the interface received illegal packets. • OFF DOWN—The interface has been shut down by OpenFlow. • Standby—The interface is a backup interface in standby state. |

New feature: Enabling PFC on all interfaces

Enabling PFC on all interfaces

About enabling PFC on all interfaces

The local device notifies the remote end to stop sending packets carrying the specified 802.1p priority if all of the following conditions exist:

- Both the local end and the remote end have PFC enabled.
- Both the local end and the remote end have the **priority-flow-control no-drop dot1p** command configured.
- The specified 802.1p priority is in the 802.1p priority list specified by the *dot1p-list* argument.
- The local end receives packets carrying the specified 802.1p priority, and the received packets cause congestion.

When congestion is eliminated, the local end notifies the remote end to continue to send packets carrying the specified 802.1p priority. In this way, the local device can forward packets carrying 802.1p priorities in the specified 802.1p priority list without packet drops.

Restrictions and guidelines

You can configure PFC for Ethernet interfaces in both system view and interface view. If you configure PFC in system view and interface view multiple times, the most recent configuration takes effect.

Procedure

1. Enter system view.
system-view
2. Enable PFC on all Ethernet interfaces.
priority-flow-control { auto | enable }
By default, PFC is disabled on all Ethernet interfaces.
3. Enable PFC for the specified 802.1p priorities on all Ethernet interfaces.
priority-flow-control no-drop dot1p dot1p-list
By default, PFC is disabled for all 802.1p priorities on all Ethernet interfaces.

Command reference

priority-flow-control (system view)

Use **priority-flow-control** to enable PFC on all Ethernet interfaces.

Use **undo priority-flow-control** to disable PFC on all Ethernet interfaces.

Syntax

```
priority-flow-control { auto | enable }  
undo priority-flow-control
```

Default

PFC is disabled on Ethernet interfaces.

Views

System view

Predefined user roles

network-admin

Parameters

auto: Specifies PFC in auto mode. In this mode, the local end automatically negotiates the PFC status with the remote end.

enable: Forcibly enables PFC.

Usage guidelines

The local device notifies the remote end to stop sending packets carrying the specified 802.1p priority if all of the following conditions exist:

- Both the local end and the remote end have PFC enabled.
- Both the local end and the remote end have the **priority-flow-control no-drop dot1p** command configured.
- The specified 802.1p priority is in the 802.1p priority list specified by the *dot1p-list* argument.
- The local end receives packets carrying the specified 802.1p priority, and the received packets cause congestion.

When congestion is eliminated, the local end notifies the remote end to continue to send packets carrying the specified 802.1p priority. In this way, the local device can forward packets carrying 802.1p priorities in the specified 802.1p priority list without packet drops.

If you execute this command in system view and Ethernet interface view multiple times, the most recent configuration takes effect.

Examples

Enable PFC on all Ethernet interfaces.

```
<Sysname> system-view  
[Sysname] priority-flow-control enable
```

Related commands

display priority-flow-control

priority-flow-control (Ethernet interface view)

priority-flow-control no-drop dot1p

Command reference

priority-flow-control no-drop dot1p (system view)

Use **priority-flow-control no-drop dot1p** to enable PFC for 802.1p priorities on all Ethernet interfaces.

Use **undo priority-flow-control no-drop dot1p** to disable PFC for 802.1p priorities on all Ethernet interfaces.

Syntax

priority-flow-control no-drop dot1p *dot1p-list*

undo priority-flow-control no-drop dot1p

Default

PFC is disabled for all 802.1p priorities on all Ethernet interfaces.

Views

System view

Predefined user roles

network-admin

Parameters

dot1p-list: Specifies an 802.1p priority (or dot1p priority) list to identify flows that are subject to PFC (for example: 1,3-5). A hyphen (-) connects two numeric values, which together indicate a continuous value range. Different values or value ranges are separated with commas (.). You can configure up to 16 characters for this argument.

Usage guidelines

The local device notifies the remote end to stop sending packets carrying the specified 802.1p priority if all of the following conditions exist:

- Both the local end and the remote end have PFC enabled.
- Both the local end and the remote end have the **priority-flow-control no-drop dot1p** command configured.
- The specified 802.1p priority is in the 802.1p priority list specified by the *dot1p-list* argument.
- The local end receives packets carrying the specified 802.1p priority, and the received packets cause congestion.

When congestion is eliminated, the local end notifies the remote end to continue to send packets carrying the specified 802.1p priority. In this way, the local device can forward packets carrying 802.1p priorities in the specified 802.1p priority list without packet drops. For more information about the 802.1p priority, see *ACL and QoS Configuration Guide*.

If you execute this command in system view and Ethernet interface view multiple times, the most recent configuration takes effect.

Examples

```
# Enable PFC for 802.1p priority 5 globally.  
<Sysname> system-view  
[Sysname] priority-flow-control no-drop dot1p 5
```

Related commands

```
display priority-flow-control  
priority-flow-control  
priority-flow-control no-drop dot1p (Ethernet interface view)
```

New feature: Configuring automatic link aggregation

Configuring automatic link aggregation

About automatic link aggregation

On a SmartMC network, you can use automatic link aggregation to aggregate the redundant physical links between devices to ease management and increase bandwidth and availability. For more information about SmartMC, see *Network Management and Monitoring Configuration Guide*.

Restrictions and guidelines

On an interface, the **port link-aggregation group** setting takes precedence over automatic link aggregation. The interface will not be added to the aggregation group of an automatically created aggregate interface if it has been the member port of a manually created aggregate interface.

For DRNI to operate correctly, do not configure automatic link aggregation on a DR system.

As a best practice to ensure correct operation of dynamic aggregation groups, do not use automatic link aggregation and automatic member port assignment together.

Prerequisites

Before you configure automatic link aggregation, enable LLDP on the peer devices.

Procedure

1. Enter system view.
system-view
2. Enable automatic link aggregation.
link-aggregation auto-aggregation enable
By default, automatic link aggregation is disabled.

Command reference

link-aggregation auto-aggregation enable

Use **link-aggregation auto-aggregation enable** to enable automatic link aggregation.

Use **undo link-aggregation auto-aggregation enable** to disable automatic link aggregation.

Syntax

link-aggregation auto-aggregation enable
undo link-aggregation auto-aggregation enable

Default

Automatic link aggregation is disabled.

Views

System view

Predefined user roles

network-admin

Usage guidelines

After you enable automatic link aggregation and LLDP on two connected devices, they automatically establish a dynamic link aggregation based on the information in incoming LLDP frames. The devices each automatically create a dynamic aggregate interface and assign the redundant ports connected to the peer to the aggregation group of that interface. If you disable LLDP or automatic link aggregation on one device, that device removes the aggregation member ports from the aggregation group, but it will retain the aggregate interface. This event will not trigger member port removal actions on the peer device.

For DRNI to operate correctly, do not use this command on a DR system.

On an interface, the **port link-aggregation group** setting takes precedence over automatic link aggregation. The interface will not be added to the aggregation group of an automatically created aggregate interface if it has been the member port of a manually created aggregate interface.

Examples

```
# Enable automatic link aggregation.  
<Sysname> system  
[Sysname] link-aggregation auto-aggregation enable
```

Related commands

- `lldp enable`
- `lldp global enable`
- `port link-aggregation group`

New feature: Disabling the default action of selecting a Selected port for dynamic aggregation groups that have not received LACPDUs

Disabling the default action of selecting a Selected port for dynamic aggregation groups that have not received LACPDUs

About the default port selection action

The default port selection action applies to dynamic aggregation groups.

This action automatically chooses the port with the lowest ID from among all up member ports as a Selected port if none of them has received LACPDUs before the LACP timeout interval expires.

After this action is disabled, a dynamic aggregation group will not have any Selected ports to forward traffic if it has not received LACPDUs before the LACP timeout interval expires.

Procedure

1. Enter system view.
`system-view`
2. Disable the default port selection action.
`lacp default-selected-port disable`

By default, the default port selection action is enabled for dynamic aggregation groups.

Command reference

lacp default-selected-port disable

Use `lacp default-selected-port disable` to disable the default port selection action for dynamic aggregation groups.

Use `undo lacp default-selected-port disable` to enable the default port selection action for dynamic aggregation groups.

Syntax

```
lacp default-selected-port disable  
undo lacp default-selected-port disable
```

Default

The default port selection action is enabled for dynamic aggregation groups.

Views

System view

Predefined user roles

network-admin

Usage guidelines

The default port selection action applies to dynamic aggregation groups.

This action automatically chooses the port with the lowest ID from among all up member ports as a Selected port if none of them has received LACPDUs before the LACP timeout interval expires.

After this action is disabled, a dynamic aggregation group will not have any Selected ports to forward traffic if it has not received LACPDUs before the LACP timeout interval expires.

Examples

```
# Disable the default port selection action.
```

```
<Sysname> system-view
```

```
[Sysname] lacp default-selected-port disable
```

New feature: Physical state change suppression on aggregate interfaces

Configuring physical state change suppression on an aggregate interface

About physical state change suppression

The physical link state of an aggregate interface is either up or down. Each time the physical link of an interface comes up or goes down, the system immediately reports the change to the CPU. The CPU then performs the following operations:

- Notifies the upper-layer protocol modules (such as routing and forwarding modules) of the change for guiding packet forwarding.
- Automatically generates traps and logs to inform users to take action.

To prevent frequent physical link flapping from affecting system performance, configure physical state change suppression. You can configure this feature to suppress link-down events, link-up events, or both. If an event of the specified type still exists when the suppression interval expires, the system reports the event to the CPU.

Restrictions and guidelines

On an interface, you can configure different suppression intervals for link-up and link-down events. If you execute the **link-delay** command multiple times for an event type, the most recent configuration takes effect.

As a best practice to reduce the impact of interface flapping on upper layer services, use this feature on the following aggregate interfaces:

- DRNI IPPs. For more information about IPPs, see *Layer 2—LAN Switching Configuration Guide*.
- Aggregate interfaces that act as cascade ports in an IRF 3.1 system. For more information about IRF 3.1, see *Virtual Technologies Configuration Guide*.

Procedure

1. Enter system view.

system-view

2. Enter aggregate interface view.

- o Enter Layer 2 aggregate interface view.

interface bridge-aggregation *interface-number*

- o Enter Layer 3 aggregate interface view.

interface route-aggregation *interface-number*

3. Configure physical state change suppression.

link-delay [**msec**] *delay-time* [**mode** { **up** | **updown** }]

By default, each time the physical link of an aggregate interface goes up or comes down, the system immediately reports the change to the CPU.

To suppress only link-down events, do not specify the **mode** keyword. To suppress only link-up events, specify the **mode** and **up** keywords. To suppress both link-down and link-up events, specify the **mode** and **updown** keywords.

Command reference

link-delay

Use **link-delay** to set the physical state change suppression interval on an aggregate interface.

Use **undo link-delay** to restore the default.

Syntax

link-delay [**msec**] *delay-time* [**mode** { **up** | **updown** }]

undo link-delay [**msec**] *delay-time* [**mode** { **up** | **updown** }]

Default

Each time the physical link of an aggregate interface goes up or comes down, the system immediately reports the change to the CPU.

Views

Layer 2 aggregate interface view

Layer 3 aggregate interface view

Predefined user roles

network-admin

Parameters

msec: Sets the physical state change suppression interval in milliseconds. If you do not specify this keyword, the suppression interval is in seconds.

delay-time: Sets the physical state change suppression interval. To report a physical state change immediately to the CPU, set the interval to 0.

- If you do not specify the **msec** keyword, the value range is 0 to 30 seconds.
- If you specify the **msec** keyword, the value range is 0 to 10000 milliseconds, and the value must be a multiple of 100.

mode up: Suppresses only the link-up events.

mode updown: Suppresses both the link-up and link-down events.

Usage guidelines

You can configure this feature to suppress link-down events, link-up events, or both. If an event of the specified type still exists when the suppression interval expires, the system reports the event.

To suppress only link-down events, execute the command without specifying the suppression mode to up or updown.

On an interface, you can configure different suppression intervals for link-up and link-down events. If you execute the **link-delay** command multiple times for an event type, the most recent configuration takes effect.

As a best practice to reduce the impact of interface flapping on upper layer services, use this command on the following aggregate interfaces:

- DRNI IPPs. For more information about IPPs, see *Layer 2—LAN Switching Configuration Guide*.
- Aggregate interfaces that act as cascade ports in an IRF 3.1 system. For more information about IRF 3.1, see *Virtual Technologies Configuration Guide*.

Examples

```
# Set the link-up and link-down event suppression interval to 500 milliseconds on Bridge-Aggregation 1.
```

```
<Sysname> system-view
[Sysname] interface bridge-aggregation 1
[Sysname-Bridge-Aggregation1] link-delay msec 500 mode updown
```

```
# Set the link-up and link-down event suppression interval to 500 milliseconds on Route-Aggregation 1.
```

```
<Sysname> system-view
[Sysname] interface route-aggregation 1
[Sysname-Route-Aggregation1] link-delay msec 500 mode updown
```

New feature: VXLAN multiport ARP entries

Configuring VXLAN multiport ARP entries

About VXLAN multiport ARP entries

A VXLAN multiport ARP entry contains an IP address, a MAC address, a VSI interface, and a VSI name. The device can use a VXLAN multiport ARP entry that has the same MAC address and VSI name as a static multicast MAC address entry or multiport unicast MAC address entry for packet forwarding. A multiport ARP entry is manually configured. It does not age out and cannot be overwritten by any dynamic ARP entry. A multiport ARP entry can overwrite a dynamic, short static or long static ARP entry. Conversely, a short static or long static ARP entry can overwrite a multiport ARP entry.

To make a multiport ARP entry effective for packet forwarding, you must configure a static multicast MAC address entry or multiport unicast MAC address entry to specify multiple output interfaces. The MAC address in a static multicast MAC address entry or multiport unicast MAC address entry can be a local or remote MAC address. A local MAC address is the MAC address of a VM in a local site. A remote MAC address is the MAC address of a VM in a remote site.

Prerequisites

A service loopback group that supports the VSI gateway ARP service must be created. The service loopback group has a minimum of one member port that is not used for any other purposes and does not have any configuration.

For a multiport ARP entry to be effective for packet forwarding, first configure a multiport unicast MAC address entry or static multicast MAC address entry with the same MAC address and VSI name as the multiport ARP entry.

Procedure

1. Enter system view.

system-view

2. Configure a multiport unicast MAC address entry or static multicast MAC address entry for VXLAN. Choose one of the following options:

- o Configure a local multiport unicast MAC address entry.

```
mac-address multiport mac-address { interface { interface-type  
interface-number service-instance instance-id }&<1-4> } vsi  
vsi-name
```

- o Configure a remote multiport unicast MAC address entry.

```
mac-address multiport mac-address { interface { tunnel  
tunnel-number1 [ to tunnel tunnel-number2 ] }&<1-4> } vsi vsi-name
```

- o Configure a local static multicast MAC address entry.

```
mac-address multicast mac-address { interface { interface-type  
interface-number service-instance instance-id }&<1-4> } vsi  
vsi-name
```

- o Configure a remote static multicast MAC address entry.

```
mac-address multicast mac-address { interface { tunnel  
tunnel-number1 [ to tunnel tunnel-number2 ] }&<1-4> } vsi vsi-name
```

3. Configure a VXLAN multiport ARP entry.

```
arp multiport ip-address mac-address vsi-interface vsi-interface-id vsi  
vsi vsi-name [ vpn-instance vpn-instance-name ]
```

Command reference

New command: arp multiport

Use **arp multiport** to configure a multiport ARP entry.

Use **undo arp** to delete an ARP entry.

Syntax

```
arp multiport ip-address mac-address vsi-interface vsi-interface-id vsi  
vsi-name [ vpn-instance vpn-instance-name ]  
undo arp ip-address [ vpn-instance-name ]
```

Default

No VXLAN multiport ARP entries exist.

Views

System view

Predefined user roles

network-admin

Parameters

ip-address: Specifies an IP address.

mac-address: Specifies a MAC address in the format of H-H-H.

vsi-interface *vsi-interface-id*: Specifies a VSI interface number in the range of 0 to 16777215.

vsi *vsi-name*: Specifies a VSI name, a case-sensitive string of 1 to 31 characters.

vpn-instance *vpn-instance-name*: Specifies an MPLS L3VPN instance to which the multiport ARP entry belongs. The *vpn-instance-name* argument represents the VPN instance name, a case-sensitive string of 1 to 31 characters. The specified VPN instance must already exist. To specify a multiport ARP entry on the public network, do not specify this option.

Usage guidelines

A service loopback group that supports the VSI gateway ARP service must be created. The service loopback group has a minimum of one member port that is not used for any other purposes and does not have any configuration.

If the VSI or VSI interface in a multiport ARP entry is deleted, the entry is also deleted.

To make the multiport ARP entry effective for packet forwarding, you must configure a static multicast MAC address entry or multiport unicast MAC address entry to specify multiple output interfaces. The MAC address entry must have the same MAC address and VSI name as the multiport ARP entry.

The VSI and VSI interface in a multiport ARP entry together identify a VXLAN IP gateway. If the VSI interface is not the gateway interface for the VSI, multiport ARP entry configuration fails. For more information about VSI interfaces and VSIs, see *VXLAN Configuration Guide*.

Examples

Configure a multiport ARP entry that contains IP address 1.1.1.1, MAC address 00e0-fc01-0000, and VSI interface 1 for VSI a.

```
<Sysname> system-view
```

```
[Sysname] arp multiport 1.1.1.1 00e0-fc01-0000 vsi-interface 1 vsi a
```

Related commands

display arp multiport

reset arp multiport

New command: mac-address multiport vsi

Use **mac-address multiport vsi** to configure a multiport unicast MAC address entry for VXLAN.

Use **undo mac-address multiport vsi** to delete a multiport unicast MAC address entry for VXLAN.

Syntax

Configure a local multiport unicast MAC address entry:

```
mac-address multiport mac-address { interface { interface-type interface-number service-instance instance-id }&<1-4> } vsi vsi-name
```

```
undo mac-address multiport [ mac-address [ interface { interface-type interface-number service-instance instance-id }&<1-4> ] ] vsi vsi-name
```

Configure a remote multiport unicast MAC address entry:

```
mac-address multiport mac-address { interface { tunnel tunnel-number1 [ to tunnel tunnel-number2 ] }&<1-4> } vsi vsi-name
```

```
undo mac-address multiport [ mac-address [ interface { tunnel tunnel-number1 [ to tunnel tunnel-number2 ] }&<1-4> ] ] vsi vsi-name
```

Default

No VXLAN multiport unicast MAC address entries exist.

Views

System view

Predefined user roles

network-admin

Parameters

mac-address: Specifies a MAC address in the format of H-H-H, excluding the all-zero MAC address. When entering a MAC address, you can omit the leading zeros in each H section. For example, enter f-e2-1 for 000f-00e2-0001.

interface { *interface-type* *interface-number* **service-instance** *instance-id* }<1-4>: Specifies up to four Ethernet service instances. The *interface-type* *interface-number* arguments represent the interface type and number. The *instance-id* argument represents an Ethernet service instance number in the range of 1 to 4096.

interface { **tunnel** *tunnel-number1* [**to tunnel** *tunnel-number2*] }<1-4>: Specifies a list of up to four VXLAN tunnel interface items. Each interface item can be an individual interface in the format of **tunnel** *tunnel-number* or a range of interfaces in the format of **tunnel** *tunnel-number1* **to tunnel** *tunnel-number2*. The value for the *tunnel-number2* argument cannot be lower than the value for the *tunnel-number1* argument. The value range for the *tunnel-number* argument is 0 to 15359.

vsi *vsi-name*: Specifies a VSI name, a case-sensitive string of 1 to 31 characters.

Usage guidelines

To send frames with a specific destination MAC address out of multiple ports, configure a multiport unicast MAC address entry. When you execute this command for the first time, the command adds an entry. When you execute the command again with the same MAC address but with different interfaces, this command adds the specified interfaces for this entry.

In VXLAN, a local MAC address is the MAC address of a VM in a local site, and a remote MAC address is the MAC address of a VM in a remote site.

To ensure successful configuration, follow these restrictions when you configure a multiport unicast MAC address entry:

- If the entry contains a local MAC address, make sure the Ethernet service instances are mapped to the VSI.
- If the entry contains a remote MAC address, make sure the tunnel interfaces are assigned to the VXLAN of the VSI.

When you use the **undo** form of this command, follow these guidelines:

- If you do not specify any parameters, this command deletes all multiport unicast MAC address entries of the specified VSI.
- If you specify a MAC address and some interfaces, this command only deletes the interfaces from the corresponding entry.

Do not configure remote MAC address entries for VXLAN tunnels that are automatically established by using EVPN. When EVPN re-establishes VXLAN tunnels, their VXLAN tunnel interface numbers might change. If you have configured remote MAC address entries, the entries are deleted when the tunnels are re-established.

You cannot configure a multiport unicast MAC address entry for the same MAC address on multiple leaf nodes or VTEPs in an EVPN network. VTEPs at a multihomed site or in a DR system do not synchronize multiport unicast MAC address entries.

Examples

Add a local multiport unicast MAC address entry that contains MAC address 000f-e201-0101 for VSI **vsi1**. Specify the output Ethernet service instances as Ethernet service instance 1 on Ten-GigabitEthernet 1/0/1 and Ten-GigabitEthernet 1/0/2.

```
<Sysname> system-view
```

```
[Sysname] mac-address multiport 000f-e201-0101 interface ten-gigabitethernet 1/0/1  
service-instance 1 ten-gigabitethernet 1/0/2 service-instance 1 vsi vsi1
```

Add a remote multiport unicast MAC address entry that contains MAC address 000f-e201-0102 for VSI **vsi1**. Specify the output interfaces as Tunnel 1, Tunnel 2, and Tunnel 3.

```
<Sysname> system-view
```

```
[Sysname] mac-address multiport 000f-e201-0102 interface tunnel 1 to tunnel 3 vsi vsi1
```

Related commands

display l2vpn mac-address (*VXLAN Command Reference*)

New command: mac-address multicast vsi

Use **mac-address multicast vsi** to configure a static multicast MAC address entry for VXLAN.

Use **undo mac-address multicast vsi** to delete a static multicast MAC address entry for VXLAN.

Syntax

Configure a local static multicast MAC address entry:

```
mac-address multicast mac-address { interface { interface-type  
interface-number service-instance instance-id }&<1-4> } vsi vsi-name
```

```
undo mac-address multicast [ mac-address [ interface { interface-type  
interface-number service-instance instance-id }&<1-4> ] ] vsi vsi-name
```

Configure a remote static multicast MAC address entry:

```
mac-address multicast mac-address { interface { tunnel tunnel-number1 [ to  
tunnel tunnel-number2 ] }&<1-4> } vsi vsi-name
```

```
undo mac-address multicast [ mac-address [ interface { tunnel  
tunnel-number1 [ to tunnel tunnel-number2 ] }&<1-4> ] ] vsi vsi-name
```

Default

No VXLAN static multicast MAC address entries exist.

Views

System view

Predefined user roles

network-admin

Parameters

mac-address: Specifies a MAC address in the format of H-H-H, excluding the all-zero MAC address. When entering a MAC address, you can omit the leading zeros in each H section. For example, enter f-e2-1 for 000f-00e2-0001.

interface { *interface-type* *interface-number* **service-instance** *instance-id* }&<1-4>: Specifies up to four Ethernet service instances. The *interface-type* *interface-number* arguments represent the interface type and number. The *instance-id* argument represents an Ethernet service instance number in the range of 1 to 4096.

interface { **tunnel** *tunnel-number1* [**to tunnel** *tunnel-number2*] }<1-4>:
Specifies a list of up to four VXLAN tunnel interface items. Each interface item can be an individual interface in the format of **tunnel** *tunnel-number* or a range of interfaces in the format of **tunnel** *tunnel-number1* **to** **tunnel** *tunnel-number2*. The value for the *tunnel-number2* argument cannot be lower than the value for the *tunnel-number1* argument. The value range for the *tunnel-number* argument is 0 to 15359.

vsi *vsi-name*: Specifies a VSI name, a case-sensitive string of 1 to 31 characters.

Usage guidelines

To send frames with a specific destination MAC address out of multiple ports, configure a static multicast MAC address entry. When you execute this command for the first time, the command adds an entry. When you execute the command again with the same MAC address and VSI name but with different interfaces, this command adds the specified interfaces for this entry.

To ensure successful configuration, follow these restrictions when you configure a static multicast MAC address entry:

- If the entry contains a local MAC address, make sure the Ethernet service instances are mapped to the VSI.
- If the entry contains a remote MAC address, make sure the tunnel interfaces are assigned to the VXLAN of the VSI.

When you use the **undo** form of this command, follow these guidelines:

- If you do not specify any parameters, this command deletes all static multicast MAC address entries of the specified VSI.
- If you specify a MAC address and some interfaces, this command only deletes the interfaces from the corresponding entry.

Do not configure remote MAC address entries for VXLAN tunnels that are automatically established by using EVPN. When EVPN re-establishes VXLAN tunnels, their VXLAN tunnel interface numbers might change. If you have configured remote MAC address entries, the entries are deleted when the tunnels are re-established.

You cannot configure a static multicast MAC address entry for the same MAC address on multiple leaf nodes or VTEPs in an EVPN network. VTEPs at a multihomed site or in a DR system do not synchronize static multicast MAC address entries.

For more information about EVPN, see *EVPN Configuration Guide*.

Examples

Add a local static multicast MAC address entry that contains MAC address 0100-5e00-0003 for VSI **vsi1**. Specify the output Ethernet service instances as Ethernet service instance 1 on Ten-GigabitEthernet 1/0/1 and Ten-GigabitEthernet 1/0/2.

```
<Sysname> system-view
```

```
[Sysname] mac-address multicast 0100-5e00-0003 interface ten-gigabitethernet 1/0/1  
service-instance 1 ten-gigabitethernet 1/0/2 service-instance 1 vsi vsi1
```

Add a remote static multicast MAC address entry that contains MAC address 0100-5e00-0003 for VSI **vsi1**. Specify the output interfaces as Tunnel 1, Tunnel 2, and Tunnel 3.

```
<Sysname> system-view
```

```
[Sysname] mac-address multicast 0100-5e00-0003 interface tunnel 1 to tunnel 3 vsi vsi1
```

Related commands

display l2vpn mac-address (*VXLAN Command Reference*)

Modified command: display arp

Syntax

```
display arp [ [ all | dynamic | multiport | static ] [ slot slot-number ]  
| vlan vlan-id | interface interface-type interface-number ] [ count |  
verbose ]
```

Views

Any view

Change description

Before modification: The **display arp** command does not display VXLAN multiport ARP entries.

Display all ARP entries.

```
<Sysname> display arp all
```

| Type: S-Static | D-Dynamic | O-Openflow | R-Rule | M-Multiport | I-Invalid |
|----------------|----------------|------------|-------------------|-------------|-----------|
| IP address | MAC address | VLAN/VSI | Interface/Link ID | Aging | Type |
| 1.1.1.1 | 02e0-f102-0023 | 1 | XGE1/0/1 | -- | S |
| 1.1.1.2 | 00e0-fc00-0001 | 12 | XGE1/0/2 | 960 | D |
| 1.1.1.3 | 00e0-fe50-6503 | 12 | Tunnell | 960 | D |
| 1.1.1.4 | 000d-88f7-9f7d | 12 | 0x1 | 960 | D |

Display detailed information about all ARP entries.

```
<Sysname> display arp all verbose
```

| | | | |
|------------|------------|-------------|------------------|
| IP address | : 1.1.1.1 | MAC address | : 02e0-f102-0023 |
| Type | : Static | Aging | : -- |
| Interface | : XGE1/0/1 | VLAN | : 1 |

VPN Instance : --

Link ID : --

VXLAN ID : --

VSI Name : --

VSI Interface : --

Nickname : 0x0000

| | | | |
|------------|------------|-------------|------------------|
| IP address | : 1.1.1.2 | MAC address | : 0015-e944-adc5 |
| Type | : Static | Aging | : 960 sec |
| Interface | : XGE1/0/2 | VLAN | : 12 |

VPN Instance : --

Link ID : --

VXLAN ID : --

VSI Name : --

VSI Interface : --

Nickname : 0x0000

| | | | |
|------------|-----------|-------------|------------------|
| IP address | : 1.1.1.3 | MAC address | : 0013-1234-0001 |
| Type | : Dynamic | Aging | : 960 sec |
| Interface | : Tunnell | VLAN | : 12 |

VPN Instance : --

Link ID : --

VXLAN ID : --

VSI Name : vpna

```

VSI Interface      : Vs11
Nickname           : 0x0000

IP address         : 1.1.1.4           MAC address      : 0012-1234-0002
Type               : Dynamic           Aging            : 960 sec
Interface          : --                VLAN             : 12
VPN Instance       : --
Link ID            : 0x1
VXLAN ID           : --
VSI Name           : vpna
VSI Interface      : Vs11
Nickname           : 0x0000

```

After modification: The **display arp** command displays VXLAN multiport ARP entries.

Display all ARP entries.

```

<Sysname> display arp all
      Type: S-Static  D-Dynamic  O-Openflow  R-Rule  M-Multiport  I-Invalid
IP address      MAC address  VLAN/VSI  Interface/Link ID  Aging Type
1.1.1.1         02e0-f102-0023 1          XGE1/0/1          --    S
1.1.1.2         00e0-fc00-0001 12         XGE1/0/2          960    D
1.1.1.3         00e0-fe50-6503 12         Tunnell1          960    D
1.1.1.4         000d-88f7-9f7d 12         0x1               960    D
1.1.1.5         00e0-fe60-5000 0          --                --    M

```

Display detailed information about all ARP entries.

```

<Sysname> display arp all verbose
IP address      : 1.1.1.1           MAC address      : 02e0-f102-0023
Type            : Static            Aging            : --
Interface       : XGE1/0/1          VLAN             : 1
VPN Instance     : --
Link ID         : --
VXLAN ID        : --
VSI Name        : --
VSI Interface    : --
Nickname        : 0x0000

IP address      : 1.1.1.2           MAC address      : 0015-e944-adc5
Type            : Static            Aging            : 960 sec
Interface       : XGE1/0/2          VLAN             : 12
VPN Instance     : --
Link ID         : --
VXLAN ID        : --
VSI Name        : --
VSI Interface    : --
Nickname        : 0x0000

IP address      : 1.1.1.3           MAC address      : 0013-1234-0001
Type            : Dynamic            Aging            : 960 sec
Interface       : Tunnell1          VLAN             : 12
VPN Instance     : --

```

```

Link ID          : --
VXLAN ID         : --
VSI Name        : vpna
VSI Interface    : Vs1l
Nickname         : 0x0000

IP address       : 1.1.1.4          MAC address      : 0012-1234-0002
Type            : Dynamic          Aging           : 960 sec
Interface       : --              VLAN             : 12
VPN Instance    : --
Link ID         : 0x1
VXLAN ID        : --
VSI Name        : vpna
VSI Interface    : Vs1l
Nickname         : 0x0000

IP address       : 1.1.1.5          MAC address      : 00e0-fe60-5000
Type            : Multiport        Aging           : --
Interface       : --              VLAN             : --
VPN instance     : --
Link ID         : --
Service instance : --
VXLAN ID        : 1
VSI name         : vpna
VSI interface    : Vs1l
Nickname         : 0x0000

```

Modified command: display l2vpn mac-address

Syntax

```
display l2vpn mac-address [ vsi vsi-name ] [ dynamic ] [ count | verbose ]
```

Views

Any view

Change description

Before modification: The **display l2vpn mac-address** command does not display multiport unicast MAC address entries or static multicast MAC address entries for VXLAN.

Display brief information about MAC address entries for all VSIs.

```
<Sysname> display l2vpn mac-address
```

```

MAC Address      State    VSI Name          Link ID/Name  Aging
0000-0000-000b   Static  vpn1              Tunnel10      NotAging
0000-0000-000c   Dynamic  vpn1              Tunnel60      Aging
0000-0000-000d   Dynamic  vpn1              Tunnel99      Aging
--- 3 mac address(es) found ---

```

Table 22 Command output

| Field | Description |
|-------|--------------|
| State | Entry state: |

| Field | Description |
|--------------|---|
| | <ul style="list-style-type: none"> • Dynamic—Local- or remote-MAC entry dynamically learned in the data plane. • Static—Static local- or remote-MAC entry. • EVPN—Remote-MAC entry advertised through EVPN. • OpenFlow—Remote-MAC entry issued by a remote controller through OpenFlow. • OVSDB—Remote-MAC entry issued by a remote controller through OVSDB. |
| Link ID/Name | <p>For a local MAC address, this field displays the interface name if the AC is a Layer 3 interface. If the AC is an Ethernet service instance, this field displays the name of the interface that hosts the Ethernet service instance for the MAC address.</p> <p>For a remote MAC address, this field displays the tunnel interface name.</p> |
| Aging | <p>Entry aging state:</p> <ul style="list-style-type: none"> • Aging. • NotAging. |

Display detailed information about MAC address entries for all VSIs.

```
<Sysname> display l2vpn mac-address verbose
MAC Address : 0000-0000-000b
VSI Name    : vpn1
VXLAN ID    : 123
Interface    : XGE1/0/1
Link ID      : 1
State        : Dynamic
Aging        : Aging
```

Table 23 Command output

| Field | Description |
|-----------|---|
| Interface | <p>For a local MAC address, this field displays the interface name if the AC is a Layer 3 interface. If the AC is an Ethernet service instance, this field displays the name of the interface that hosts the Ethernet service instance for the MAC address.</p> <p>For a remote MAC address, this field displays the tunnel interface name.</p> |
| Link ID | Link ID that uniquely identifies an AC or a VXLAN tunnel on a VSI. |
| State | <p>Entry state:</p> <ul style="list-style-type: none"> • Dynamic—Local- or remote-MAC entry dynamically learned in the data plane. • Static—Static local- or remote-MAC entry. • EVPN—Remote-MAC entry advertised through EVPN. • OpenFlow—Remote-MAC entry issued by a remote controller through OpenFlow. • OVSDB—Remote-MAC entry issued by a remote controller through OVSDB. |
| Aging | <p>Entry aging state:</p> <ul style="list-style-type: none"> • Aging. • NotAging. |

After modification: The **display l2vpn mac-address** command displays multiport unicast MAC address entries and static multicast MAC address entries for VXLAN.

Display brief information about MAC address entries for all VSIs.

```
<Sysname> display l2vpn mac-address
```

| MAC Address | State | VSI Name | Link ID/Name | Aging |
|----------------|------------------|----------|--------------|----------|
| 0000-0000-000b | Static | vpn1 | Tunnel10 | NotAging |
| 0000-0000-000c | Dynamic | vpn1 | Tunnel60 | Aging |
| 0000-0000-000d | Dynamic | vpn1 | Tunnel99 | Aging |
| 0100-5e00-0003 | Multicast | vpnb | XGE1/0/3 | NotAging |

--- 4 mac address(es) found ---

Table 24 Command output

| Field | Description |
|--------------|---|
| State | <p>Entry state:</p> <ul style="list-style-type: none"> Dynamic—Local- or remote-MAC entry dynamically learned in the data plane. Static—Static local- or remote-MAC entry. EVPN—Remote-MAC entry advertised through EVPN. OpenFlow—Remote-MAC entry issued by a remote controller through OpenFlow. OVSDb—Remote-MAC entry issued by a remote controller through OVSDb. Multiport—Remote or local multipoint unicast MAC address entry. Multicast—Remote or local static multicast MAC address entry. |
| Link ID/Name | <p>For a local MAC address, this field displays the interface name if the AC is a Layer 3 interface. If the AC is an Ethernet service instance, this field displays the name of the interface that hosts the Ethernet service instance for the MAC address.</p> <p>For a remote MAC address, this field displays the tunnel interface name.</p> |
| Aging | <p>Entry aging state:</p> <ul style="list-style-type: none"> Aging. NotAging. |

Display detailed information about MAC address entries for all VSIs.

```
<Sysname> display l2vpn mac-address verbose
```

```
MAC Address : 0000-0000-000b
```

```
VSI Name : vpn1
```

```
VXLAN ID : 123
```

```
Interface : XGE1/0/1
```

```
Link ID : 1
```

```
State : Dynamic
```

```
Aging : Aging
```

```
MAC Address : 0002-0003-0004
```

```
VSI Name : vpna
```

```
VXLAN ID : 10
```

```
State : Multiport
```

```
Aging : NotAging
```

```
Interface Link ID
```

```
GE1/0/2 0x0
```

```
Tunnel8 0x5000002
```

```
Tunnel9 0x5000003
```

```

MAC Address : 0100-5e00-0003
VSI Name    : vpnb
VXLAN ID    : 20
State       : ☒
Aging       : NotAging
Interface    Link ID
GE1/0/3      0x0
Tunnel12     0x5000004

```

Table 25 Command output

| Field | Description |
|-----------|---|
| Interface | For a local MAC address, this field displays the interface name if the AC is a Layer 3 interface. If the AC is an Ethernet service instance, this field displays the name of the interface that hosts the Ethernet service instance for the MAC address. For a remote MAC address, this field displays the tunnel interface name. |
| Link ID | Link ID that uniquely identifies an AC or a VXLAN tunnel on a VSI. |
| State | Entry state: <ul style="list-style-type: none"> • Dynamic—Local- or remote-MAC entry dynamically learned in the data plane. • Static—Static local- or remote-MAC entry. • EVPN—Remote-MAC entry advertised through EVPN. • OpenFlow—Remote-MAC entry issued by a remote controller through OpenFlow. • OVSDb—Remote-MAC entry issued by a remote controller through OVSDb. • Multiport—Remote or local multiport unicast MAC address entry. • Multicast—Remote or local static multicast MAC address entry. |
| Aging | Entry aging state: <ul style="list-style-type: none"> • Aging. • NotAging. |

New feature: Enabling recording user port move for ARP

Enabling recording user port move for ARP

About recording user port move for ARP

This feature enables the device to detect and record user port move events. A user port move occurs if an incoming NA packet has the same source IP address and source MAC address as an existing ARP entry but a different port. The device generates a user port move record, logs the move event, and sends the log to the information center. For information about the log destination and output rule configuration in the information center, see the information center in *Network Management and Monitoring Configuration Guide*.

When a large number of users change their access ports, the device will generate a large number of user port move records, degrading device performance. As a best practice, disable recording user port move in such situation.

Procedure

1. Enter system view.
system-view
2. Enable recording user port move for ARP.
arp user-move record enable
By default, recording user port move for ARP is disabled.

Display and maintenance commands for ARP

Execute the **display** command in any view.

| Task | Command |
|--|---|
| Display user port move records in ARP. | display arp user-move record [<i>slot slot-number</i>] |

Command reference

arp user-move record enable

Use **arp user-move record enable** to enable recording user port moves for ARP.

Use **undo arp user-move record enable** to disable recording user port moves for ARP.

Syntax

```
arp user-move record enable
undo arp user-move record enable
```

Default

Recording user port moves for ARP is disabled.

Views

System view

Predefined user roles

network-admin

Usage guidelines

An IRF member device can generate a maximum of 10 user port move logs per second.

To display more user port move records, use the **display arp user-move record** command.

Examples

```
# Enabling recording user port move.
<Sysname> system-view
[Sysname] arp user-move record enable
```

Related commands

```
display arp user-move record
```

display arp user-move record

Use **display arp user-move record** to display user port move records in ARP.

Syntax

```
display arp user-move record [ slot slot-number ]
```

Views

Any view

Predefined user roles

network-admin
network-operator

Parameters

slot *slot-number*: Specifies an IRF member device by its member ID. If you do not specify a member device, this command displays user port move records for all member devices.

Usage guidelines

An IRF member device can save a maximum of 200 user port move records. When the number of saved user port move records reaches the upper limit, new records overwrite old ones.

Examples

Display all user port move records for ARP.

```
<Sysname> display arp user-move record
```

```
IP address: 10.1.1.1
```

```
MAC address: 0001-0201-0e81
```

```
System time: 2018-02-02 11:22:29
```

```
Move count: 1
```

```
Log suppress count: 0
```

```
Before:
```

```
    interface: Ten-GigabitEthernet1/0/1
```

```
    SVLAN/CVLAN: 100/2
```

```
After:
```

```
    interface: Ten-GigabitEthernet1/0/2
```

```
    SVLAN/CVLAN: 100/2
```

```
IP address: 10.1.1.2
```

```
MAC address: 0001-0201-0e82
```

```
System time: 2018-02-02 10:20:30
```

```
Move count: 1
```

```
Log suppress count: 0
```

```
Before:
```

```
    interface: Ten-GigabitEthernet1/0/1
```

```
    SVLAN/CVLAN: 100/--
```

```
After:
```

```
    interface: Ten-GigabitEthernet1/0/2
```

```
    SVLAN/CVLAN: 100/--
```

Table 26 Command output

| Field | Description |
|-------------|--|
| IP address | IP address of the user. |
| MAC address | MAC address of the user. |
| System time | Time when the user port move occurred. |

| Field | Description |
|--------------------|---|
| Move count | Number of times the user port moved. |
| Log suppress count | Number of times user port move logs are suppressed. |
| Before: | Information before the user port move. |
| interface | Interface information in the ND entry. |
| SVLAN/CVLAN | ID of the outer VLAN or inner VLAN in the ND entry. This field displays hyphens (--) if the ND entry does not belong to any outer VLAN or inner VLAN. |
| After: | Information after the user port move. |

New feature: Enabling recording user IP address conflicts

Enabling recording user IP address conflicts

About recording user IP address conflicts

This feature enables the device to detect and record user IP address conflicts. The device determines that a conflict occurs if an incoming non-gratuitous ARP packet has the same sender IP address as an existing ARP entry but a different sender MAC address. The device generates a user IP address conflict record, logs the conflict, and sends the log to the information center. For information about the log destination and output rule configuration in the information center, see the information center in *Network Management and Monitoring Configuration Guide*.

Procedure

1. Enter system view.
system-view
2. Enable recording user IP address conflicts.
arp user-ip-conflict record enable
By default, recording user IP address conflicts is disabled.

Display and maintenance commands for ARP

Execute display commands in any view and reset commands in user view.

Display user IP address conflicts.

```
display arp user-ip-conflict record [ slot  
slot-number ]
```

Command reference

arp user-ip-conflict record enable

Use **arp user-ip-conflict record enable** to enable recording user IP address conflicts.

Use **undo arp user-ip-conflict record enable** to disable recording user IP address conflicts.

Syntax

```
arp user-ip-conflict record enable
undo arp user-ip-conflict record enable
```

Default

Recording user IP address conflicts is disabled.

Views

System view

Predefined user roles

network-admin

Usage guidelines

This feature enables the device to detect and record user IP address conflicts. The device determines that a conflict occurs if an incoming non-gratuitous ARP packet has the same sender IP address as an existing ARP entry but a different sender MAC address. The device generates a user IP address conflict record, logs the conflict, and sends the log to the information center. For information about the log destination and output rule configuration, see the information center in *Network Management and Monitoring Configuration Guide*.

An IRF member device can generate a maximum of 10 user IP address conflict logs per second.

To display user IP address conflict records, use the **display arp user-ip-conflict record** command.

Examples

```
# Enable recording user IP address conflicts.
<Sysname> system-view
[Sysname] arp user-ip-conflict record enable
```

Related commands

```
display arp user-ip-conflict record
```

display arp user-ip-conflict record

Use **display arp user-ip-conflict record** to display user IP address conflict records.

Syntax

```
display arp user-ip-conflict record [ slot slot-number ]
```

Views

Any view

Predefined user roles

network-admin
network-operator

Parameters

slot slot-number: Specifies an IRF member device by its member ID. If you do not specify a member device, this command displays user IP address conflict records for the master device.

Usage guidelines

Each IRF member device can save a maximum of 200 user IP address conflict records.

If the maximum number is reached, a new record will override the earliest record.

Examples

```
# Display all user IP address conflict records.
<Sysname> display arp user-ip-conflict record
IP address: 10.1.1.1
System time: 2018-02-02 11:22:29
Conflict count: 1
Log suppress count: 0
Old interface: Ten-GigabitEthernet1/0/1
New interface: Ten-GigabitEthernet1/0/2
Old SVLAN/CVLAN: 100/2
New SVLAN/CVLAN: 100/2
Old MAC: 00e0-ca63-8141
New MAC: 00e0-ca63-8142

IP address: 10.1.1.2
System time: 2018-02-02 10:20:30
Conflict count: 1
Log suppress count: 0
Old interface: Ten-GigabitEthernet1/0/1
New interface: Ten-GigabitEthernet1/0/2
Old SVLAN/CVLAN: 100/--
New SVLAN/CVLAN: 100/--
Old MAC: 00e0-ca63-8141
New MAC: 00e0-ca63-8142
```

Table 27 Command output

| Field | Description |
|--------------------|---|
| IP address | IP address of a user. |
| System time | Time when the user IP address conflict occurred. |
| Conflict count | Number of times that conflicts for the IP address. |
| Log suppress count | Number of times that user IP address conflict logs are suppressed. |
| Old interface | Output interface in the old ARP entry. |
| New interface | Output interface in the new ARP entry. |
| Old SVLAN/CVLAN | ID of the outer VLAN or inner VLAN in the old ARP entry. This field displays hyphens (--) if the ARP entry does not belong to any outer VLAN or inner VLAN. |
| New SVLAN/CVLAN | ID of the outer VLAN or inner VLAN in the new ARP entry. This field displays hyphens (--) if the ARP entry does not belong to any outer VLAN or inner VLAN. |
| Old MAC | MAC address in the old ARP entry. |
| New MAC | MAC address in the new ARP entry. |

Related commands

```
arp user-ip-conflict record enable
```

New feature: Enabling interface consistency check for ARP and MAC address entries

Enabling interface consistency check for ARP and MAC address entries

About interface consistency check for ARP and MAC address entries

In an instable network, the receiving interface for packets from a user might change. The interface in the MAC address entry can be updated immediately while the interface in the ARP entry cannot. In this case, the packets matching the ARP entry will be sent out of an incorrect interface. To solve this problem, you can use this feature to periodically check the interface consistency for the ARP and MAC address entry of a user. If the interfaces are not the same, ARP sends ARP requests in the VLAN of the ARP entry and updates the entry with the ARP reply receiving interface.

Procedure

1. Enter system view.

```
system-view
```

2. Enabling interface consistency check for ARP and MAC address entries.

```
arp mac-interface-consistency check enable
```

By default, enabling interface consistency check for ARP and MAC address entries is disabled.

Command reference

arp mac-interface-consistency check enable

Use **arp mac-interface-consistency check enable** to enable interface consistency check for ARP and MAC address entries.

Use **undo arp mac-interface-consistency check enable** to disable this feature.

Syntax

```
arp mac-interface-consistency check enable
```

```
undo arp mac-interface-consistency check enable
```

Default

Enabling interface consistency check for ARP and MAC address entries is disabled.

Views

System view

Predefined user roles

network-admin

Usage guidelines

In an instable network, the receiving interface for packets from a user might change. The interface in the MAC address entry can be updated immediately while the interface in the ARP entry cannot. In this case, the packets matching the ARP entry will be sent out of an incorrect interface. To solve this problem, you can use this feature to periodically check the interface consistency for the ARP and MAC address entry of a user. If the interfaces are not the same, ARP sends ARP requests in the VLAN of the ARP entry and updates the entry with the ARP reply receiving interface.

Use **display mac-address** to display MAC address entries.

Examples

```
# Enable interface consistency check for ARP and MAC address entries.
<Sysname> system-view
[Sysname] arp mac-interface-consistency check enable
```

New feature: Specifying a DHCP relay address pool for DHCP clients

Specifying a DHCP relay address pool for DHCP clients

About specifying a DHCP relay address pool for DHCP clients

After you configure multiple DHCP relay address pools on a DHCP relay agent, you can specify these pools on an interface. To match DHCP clients based on options, you can define option settings when you specify the relay address pools.

If you specify multiple DHCP relay address pools on an interface, the relay agent selects a DHCP relay address pool for a DHCP client as follows:

1. Compares option values in the DHCP request in descending order against option values in DHCP relay address pools.
 - If a match (other than 60) is found, the matching process stops and the relay agent selects that matching relay address pool.
 - If the matching option value is 60, the relay agent continues to compare the Option 60 content in the request and the Option 60 string in the relay address pool:
 - If the Option 60 content matches the string, the relay address pool is selected.
 - If the Option 60 content does not match the string, the relay address pool is not selected. If another relay address pool is specified to match Option 60 but has no Option 60 string defined, the relay agent selects that relay address pool.
2. If still no DHCP relay address pool is matched, the relay agent selects the DHCP relay address pool with no options specified.

Restrictions and guidelines

If you specify DHCP servers by configuring both of the following methods on an interface, the DHCP relay address pool setting takes effect.

- Specify DHCP relay address pools by using the **dhcp relay pool** command.
- Specify DHCP servers directly on an interface by using the **dhcp relay server-address** command.

When you specify a DHCP relay address pool on an interface to define the DHCP servers, make sure the **remote-server** command is configured in the DHCP relay address pool. Otherwise, the relay agent drops DHCP requests. The DHCP requests are not forwarded to any DHCP server even if the **dhcp relay server-address** command is configured.

Procedure

1. Enter system view.
system-view
2. Create a DHCP relay address pool and enter its view.
dhcp server ip-pool pool-name
By default, no DHCP relay address pools exist.
3. Specify DHCP servers in the DHCP relay address pool.

```
remote-server ip-address&<1-8> [ public | vpn-instance vpn-instance-name ]
```

By default, no DHCP server is specified in the DHCP relay address pool.

4. Specify gateway addresses for the clients matching the DHCP relay address pool.

```
gateway-list ip-address&<1-64> export-route
```

By default, no gateway address is specified.

5. Specify the DHCP server selecting algorithm.

```
remote-server algorithm { master-backup | polling }
```

By default, the **polling** algorithm is used. The DHCP relay agent forwards DHCP requests to all DHCP servers at the same time.

6. Return to system view.

```
quit
```

7. Enter interface view.

```
interface interface-type interface-number
```

8. Specify a DHCP relay address pool for DHCP clients.

```
dhcp relay pool pool-name [ option { 60 [ option-text ] | code } ]
```

Command reference

Modified command:remote-server

Old syntax

```
remote-server ip-address&<1-8>  
undo remote-server [ ip-address&<1-8> ]
```

New syntax

```
remote-server ip-address&<1-8> [ public | vpn-instance vpn-instance-name ]  
undo remote-server [ ip-address&<1-8> ]
```

Views

DHCP address pool view

Parameters

ip-address&<1-8>: Specifies a space-separated list of up to eight DHCP server addresses.

public: Specifies the DHCP servers on the public network.

vpn-instance *vpn-instance-name*: Specifies the name of the MPLS L3VPN instance to which the DHCP servers belong. The instance name is a case-sensitive string of 1 to 31 characters.

dhcp relay pool

Use **dhcp relay pool** to specify a DHCP relay address pool for DHCP clients.

Use **undo dhcp relay pool** to remove a DHCP relay address pool.

Syntax

```
dhcp relay pool pool-name [ option { 60 [ option-text ] | code } ]  
undo dhcp relay pool [ pool-name [ option { 60 [ option-text ] | code } ] ]
```

Default

No DHCP relay address pool is specified for DHCP clients.

Views

Layer 3 Ethernet interface view

Layer 3 Ethernet subinterface view

Layer 3 aggregate interface view

Layer 3 aggregate subinterface view

Predefined user roles

network-admin

Parameters

pool-name: Specifies a DHCP relay address pool by its name, a case-insensitive string of 1 to 63 characters.

option { **60** [*option-text*] | *code* }: Specifies the option setting to match DHCP requests. If you do not specify this option, the DHCP relay address pool can match all DHCP requests.

- The **60** keyword specifies Option 60. The *option-text* argument specifies Option 60 content, a case-sensitive string of 1 to 128 characters. This argument is used to match Option 60 content in DHCP requests and is configurable only when the **60** keyword is specified.
- The *code* argument specifies an option value in the range of 1 to 59 or 61 to 254.

Usage guidelines

After you configure multiple DHCP relay address pools on a DHCP relay agent, you can specify these pools on an interface. To match DHCP clients based on options, you can define option settings when you specify the relay address pools.

If you specify multiple DHCP relay address pools on an interface, the relay agent selects a DHCP relay address pool for a DHCP client as follows:

1. Compares option values in the DHCP request in descending order against option values in DHCP relay address pools.
 - If a match (other than 60) is found, the matching process stops and the relay agent selects that matching relay address pool.
 - If the matching option value is 60, the relay agent continues to compare the Option 60 content in the request and the Option 60 string in the relay address pool:
 - If the Option 60 content matches the string, the relay address pool is selected.
 - If the Option 60 content does not match the string, the relay address pool is not selected. If another relay address pool is specified to match a request carrying Option 60 but has no Option 60 string defined, the relay agent selects that relay address pool.
2. If still no DHCP relay address pool is matched, the relay agent selects the DHCP relay address pool with no options specified.

If you specify DHCP servers by configuring both the following methods on an interface, the DHCP relay address pool setting takes effect.

- Specify DHCP relay address pools by using the **dhcp relay pool** command.
- Specify DHCP servers directly on an interface by using the **dhcp relay server-address** command.

When you specify a DHCP relay address pool on an interface to define the DHCP servers, make sure the **remote-server** command is configured in the DHCP relay address pool. Otherwise, the relay agent drops DHCP requests. The DHCP requests are not forwarded to any DHCP server even if the **dhcp relay server-address** command is configured.

If you execute this command multiple times to specify different DHCP address pools for the same DHCP client, the most recent configuration takes effect.

If you do not specify a DHCP address pool, the **undo dhcp relay pool** command removes all DHCP address pools specified on the interface .

Examples

Specify DHCP relay address pool **pool1** on VLAN-interface 2.

```
<Sysname> system-view
[Sysname] interface vlan-interface 2
[Sysname-Vlan-interface2] dhcp relay pool pool1
```

New feature: Configuring the padding mode for the Vendor-Specific sub-option

Configuring the padding mode for the Vendor-Specific sub-option

Restrictions and guidelines

- The Option 82 configuration on a Layer 2 Ethernet interface that has been added to an aggregation group does not take effect unless the interface leaves the aggregation group.
- To support Option 82, you must configure Option 82 on both the DHCP server and the DHCP snooping device.
- If Option 82 contains the device name, the device name must contain no spaces. Otherwise, DHCP snooping drops the message. You can use the **sysname** command to specify the device name. For more information about this command, see *Fundamentals Command Reference*.
- DHCP snooping uses "outer VLAN tag.inner VLAN tag" to fill the VLAN ID field of sub-option 1 in verbose padding format if either of the following conditions exists:
 - DHCP snooping and QinQ work together.
 - DHCP snooping receives a DHCP packet with two VLAN tags.

For example, if the outer VLAN tag is 10 and the inner VLAN tag is 20, the VLAN ID field is 000a.0014. The hexadecimal digit **a** represents the outer VLAN tag 10, and the hexadecimal digit **14** represents the inner VLAN tag 20.

Procedure

1. Enter system view.
system-view
2. Enter interface view.
interface *interface-type interface-number*
3. Enable DHCP snooping to support Option 82.
dhcp snooping information enable
By default, DHCP snooping does not support Option 82.
4. Configure the padding mode for the Vendor-Specific sub-option.
dhcp snooping information vendor-specific [**vlan** *vlan-id*] **bas**
[**node-identifier** { **mac** | **sysname** | **user-defined** *string* }]
By default, the device does not pad the Vendor-Specific sub-option.

Command reference

dhcp snooping information vendor-specific

Use **dhcp snooping information vendor-specific** to configure the padding mode for the Vendor-Specific sub-option.

Use **undo dhcp snooping information vendor-specific** to restore the default.

Syntax

```
dhcp snooping information vendor-specific [ vlan vlan-id ] bas  
[ node-identifier { mac | sysname | user-defined string } ]  
undo dhcp snooping information vendor-specific [ vlan vlan-id ]
```

Default

The device does not pad the Vendor-Specific sub-option.

Views

Layer 2 Ethernet interface view

Layer 2 aggregate interface view

Predefined user roles

network-admin

Parameters

vlan *vlan-id*: Pads the Vendor-Specific sub-option for packets received from the specified VLAN. If you do not specify a VLAN, the device pads the Vendor-Specific sub-option for all packets received on the interface.

bas: Specifies the bas mode to pad the Vendor-Specific sub-option.

node-identifier: Specifies the access node identifier. If you do not specify this keyword, the device pads the Vendor-Specific sub-option with the bridge MAC address of the access node as the node identifier. The padding format for the Vendor-Specific sub-option is ASCII.

- **mac**: Uses the bridge MAC address of the access node as the node identifier.
- **sysname**: Uses the device name as the node identifier. You can set the device name by using the **sysname** command in system view. If the **sysname** keyword is specified, make sure the device name does not include any spaces. Otherwise, the DHCP snooping device fails to add the Vendor-Specific sub-option. If the device name contains more than 50 characters, only the first 50 characters are padded.
- **user-defined *string***: Uses a case-sensitive string of 1 to 50 characters as the node identifier. Do not include any spaces in the string.

Usage guidelines

After you configure this command, the DHCP snooping device pads the Vendor-Specific sub-option after receiving a DHCP request. The device forwards the DHCP request without padding the Vendor-Specific sub-option if the following conditions exist:

- The **dhcp snooping information strategy append** command is configured.
- The length of Option 82 in the request reaches the upper limit.

Examples

Pad the Vendor-Specific sub-option in bas mode with the device name as the node identifier.

```
<Sysname> system-view
```

```
[Sysname] interface ten-gigabitethernet 1/0/1
```

```
[Sysname-Ten-GigabitEthernet1/0/1] dhcp snooping information enable
[Sysname-Ten-GigabitEthernet1/0/1] dhcp snooping information vendor-specific bas
node-identifier sysname
```

New feature: Configuring IPv6 ND direct route advertisement

Configuring IPv6 ND direct route advertisement

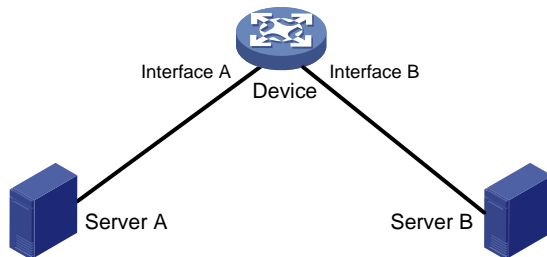
About IPv6 ND direct route advertisement

The ND direct route advertisement feature advertises host routes instead of advertising the network route.

Layer 3 access network application

As shown in [Figure 2](#), Server A and Server B access the device through Layer 3 interfaces Interface A and Interface B, respectively. With ND direct route advertisement feature enabled on Interface A and Interface B, the device generates host routes to Server A and Server B, which saves bandwidth in route advertisement.

Figure 2 Layer 3 access network application



Procedure

1. Enter system view.
system-view
2. Enter interface view.
interface *interface-type* *interface-number*
3. Enable ND direct route advertisement.
ipv6 nd route-direct advertise
By default, ND direct route advertisement is disabled.

Command reference

ipv6 nd route-direct advertise

Use **ipv6 nd route-direct advertise** to enable ND direct route advertisement.

Use **undo ipv6 nd route-direct advertise** to disable ND direct route advertisement.

Syntax

```
ipv6 nd route-direct advertise
undo ipv6 nd route-direct advertise
```

Default

The ND direct route advertisement feature is disabled.

Views

Interface view

Predefined user roles

network-admin

Examples

```
# Enable ND direct route advertisement for VLAN-interface 100.
<Sysname> system-view
[Sysname] interface vlan-interface 100
[Sysname-Vlan-interface100] ipv6 nd route-direct advertise

# Enable ND direct route advertisement for VSI-interface 1.
<Sysname> system-view
[Sysname] interface vsi-interface 1
[Sysname-Vsi-interface1] ipv6 nd route-direct advertise
```

New feature: Configuring ND snooping in a VXLAN

Configuring ND snooping in a VXLAN

About ND snooping in a VXLAN

The ND snooping feature in a VXLAN learns the source MAC addresses, source IPv6 addresses, VSI name, and link IDs to build the ND snooping table. For more information about VXLAN, see VXLAN overview in *VXLAN Configuration Guide*.

ND snooping entries can be used by ND detection to prevent spoofing attacks. For more information about the ND attack detection, see *Security Configuration Guide*.

ND snooping entries can be used by IPv6 source guard to prevent spoofing attacks. For more information about IPv6 source guard, see *Security Configuration Guide*.

After ND snooping is enabled in a VSI, the device uses the following mechanisms to create and maintain ND snooping entries.

Creation of ND snooping entries

With this feature enabled in a VXLAN, all receiving ND messages are delivered to the CPU. The CPU obtains the source IPv6 addresses, source MAC addresses, VSI name, and link IDs to create ND snooping entries.

Aging of ND snooping entries

The aging time of an ND snooping entry is 5 minutes. If no matching packet is received within 30 seconds before the end of the entry aging time, the entry is set to TENTATIVE status. At the same time, the device sends a NS message for the entry. If the source IP and source MAC in the received ND message matches the ND snooping entry, the device sets the entry to VALID and resets the aging time.

When the aging time of an ND snooping entry is reached, the entry is deleted.

Procedure

1. Enter system view.
system-view
 2. Create a VSI and enter its view.
vsi *vsi-name*
 3. Enable ND snooping for IPv6 addresses. Choose the options to configure as needed:
 - o Enable ND snooping for global unicast addresses.
ipv6 nd snooping enable global
 - o Enable ND snooping for link-local addresses.
ipv6 nd snooping enable link-local
- By default, ND snooping is disabled for IPv6 global unicast addresses and link-local addresses.

Command reference

Modified command: display ipv6 nd snooping count

Old syntax

```
display ipv6 nd snooping count [ interface interface-type  
interface-number ]
```

New syntax

```
display ipv6 nd snooping count vlan [ interface interface-type  
interface-number ]
```

Views

Any view

Change description

The **vlan** keyword was added to the **display ipv6 nd snooping** command to display the number of IPv6 ND snooping entries in VLANs.

Modified command: ipv6 nd snooping enable global

Syntax

```
ipv6 nd snooping enable global  
undo ipv6 nd snooping enable global
```

Default

ND snooping is disabled for global unicast addresses.

Views

VLAN view

VSI view

Change description

Before modification: The **ipv6 nd snooping enable global** command is not supported in VSI view.

After modification: The `ipv6 nd snooping enable global` command is supported in VSI view.

Modified command: `ipv6 nd snooping enable link-local`

Syntax

```
ipv6 nd snooping enable link-local
undo ipv6 nd snooping enable link-local
```

Default

ND snooping is disabled for link-local addresses.

Views

VLAN view
VSI view

Change description

Before modification: The `ipv6 nd snooping enable link-local` command is not supported in VSI view.

After modification: The `ipv6 nd snooping enable link-local` command is supported in VSI view.

New command: `display ipv6 nd snooping count vsi`

Use `display ipv6 nd snooping count vsi` to display the number of IPv6 ND snooping entries in the specified VSI.

Syntax

```
display ipv6 nd snooping count vsi [ vsi-name ]
```

Views

Any view

Predefined user roles

network-admin
network-operator

Parameters

vsi-name: Specifies a VSI name, a case-sensitive string of 1 to 31 characters. If you do not specify a VSI, this command displays ND snooping entries for all VSIs.

Examples

Display the total number of IPv6 ND snooping entries in all VSIs.

```
<Sysname> display ipv6 nd snooping count vsi
Total entries for VSIs: 5
```

Display the total number of IPv6 ND snooping entries in VSI **vsi1**.

```
<Sysname> display ipv6 nd snooping count vsi vsi1
Total entries for vsi1: 2
```

Table 28 Command output

| Field | Description |
|------------------------|---|
| Total entries for VSIs | Total number of IPv6 ND snooping entries in all VSIs. |

| Field | Description |
|------------------------|---|
| Total entries for vsi1 | Total number of IPv6 ND snooping entries in VSI vsi1 . |

Related commands

```

ipv6 nd snooping enable global
ipv6 nd snooping enable link-local
reset ipv6 nd snooping vsi

```

New command: display ipv6 nd snooping vsi

Use **display ipv6 nd snooping vsi** to display ND snooping entries in the specified VSI.

Syntax

```
display ipv6 nd snooping vsi [ vsi-name ] [ slot slot-number ]
```

Views

Any view

Predefined user roles

```

network-admin
network-operator

```

Parameters

vsi-name: Specifies a VSI name, a case-sensitive string of 1 to 31 characters. If you do not specify a VSI, this command displays ND snooping entries for all VSIs.

slot slot-number: Specifies an IRF member device by its member ID. If you do not specify a member device, this command displays ND snooping entries for the master device.

Examples

Display ND snooping entries in VSI **vsi1**.

```
<Sysname> display ipv6 nd snooping vsi vsi1
```

| | | | | |
|--------------|----------------|----------|---------|------------|
| IPv6 address | MAC address | VSI name | Link ID | Aging(min) |
| 1::2 | 0000-1234-0c01 | vsi1 | 0x70000 | 5 |

Table 29 Command output

| Field | Description |
|--------------|--|
| IPv6 address | IPv6 address in the ND snooping entry. |
| MAC address | MAC address in the ND snooping entry. |
| VSI name | Name of the VSI to which the ND snooping entry belongs. |
| Link ID | Link ID that uniquely identifies an AC or a VXLAN tunnel on a VSI. |
| Aging(min) | Remaining aging time of the ND snooping entry, in minutes. |

Related commands

```

ipv6 nd snooping enable global
ipv6 nd snooping enable link-local
reset ipv6 nd snooping vsi

```

New command: reset ipv6 nd snooping vsi

Use **reset ipv6 nd snooping vsi** to clear ND snooping entries in VSIs.

Syntax

```
reset ipv6 nd snooping vsi [ vsi-name ]
```

Views

User view

Predefined user roles

network-admin

Parameters

vsi-name: Specifies a VSI name, a case-sensitive string of 1 to 31 characters. If you do not specify a VSI, this command clear ND snooping entries in all VSIs.

Examples

```
# Clear ND snooping entries in VSI vsi1.  
<Sysname> reset ipv6 nd snooping vsi vsi1
```

Related commands

```
display ipv6 nd snooping count vsi  
display ipv6 nd snooping vsi
```

New feature: Enabling recording user IPv6 address conflicts

Enabling recording user IPv6 address conflicts

About recording user IPv6 address conflicts

This feature enables the device to detect and record user IPv6 address conflicts. A conflict occurs if an incoming NA packet has the same source IP address as an existing ND entry but a different source MAC address. The device generates a user IPv6 address conflict record, logs the conflict, and sends the log to the information center. For information about the log destination and output rule configuration in the information center, see the information center in *Network Management and Monitoring Configuration Guide*.

Procedure

1. Enter system view.
system-view
2. Enable recording user IPv6 address conflicts.
ipv6 nd user-ip-conflict record enable
By default, recording user IPv6 address conflicts is disabled.

Command reference

ipv6 nd user-ip-conflict record enable

Use **ipv6 nd user-ip-conflict record enable** to enable recording user IPv6 address conflicts.

Use **undo ipv6 nd user-ip-conflict record enable** to disable recording user IPv6 address conflicts.

Syntax

```
ipv6 nd user-ip-conflict record enable
undo ipv6 nd user-ip-conflict record enable
```

Default

Recording user IPv6 address conflicts is disabled.

Views

System view

Predefined user roles

network-admin

Usage guidelines

This feature enables the device to detect and record user IPv6 address conflicts. A conflict occurs if an incoming NA packet has the same source IPv6 address as an existing ND entry but a different source MAC address. The device generates a user IPv6 address conflict record, logs the conflict, and sends the log to the information center. For information about the log destination and output rule configuration in the information center, see the information center in *Network Management and Monitoring Configuration Guide*.

An IRF member device can generate a maximum of 10 user IPv6 address conflict logs per second. The device records times user IPv6 address conflict log generation is suppressed after the related log count reaches the upper limit. An IRF member device can save a maximum of 200 user IPv6 address conflict records.

When the number of saved user IPv6 address conflict records reaches the upper limit, new records overwrite old ones.

Examples

```
# Enable recording user IPv6 address conflicts.
<Sysname> system-view
[Sysname] ipv6 nd user-ip-conflict record enable
```

Related commands

```
display ipv6 nd user-ip-conflict record
```

display ipv6 nd user-ip-conflict record

Use **display ipv6 nd user-ip-conflict record** to display user IPv6 address conflict records.

Syntax

```
display ipv6 nd user-ip-conflict record [ slot slot-number ]
```

Views

Any view

Predefined user roles

network-admin
network-operator

Parameters

slot *slot-number*: Specifies an IRF member device by its member ID. If you do not specify a member device, this command displays user IP address conflict records for all member devices.

Examples

Display all user IPv6 address conflict records.

```
<Sysname> display ipv6 nd user-ip-conflict record
```

```
IPv6 address: 10::1
```

```
System time: 2018-02-02 11:22:29
```

```
Conflict count: 1
```

```
Log suppress count: 0
```

```
Old interface: GigabitEthernet1/0/1
```

```
New interface: GigabitEthernet1/0/2
```

```
Old SVLAN/CVLAN: 100/2
```

```
New SVLAN/CVLAN: 100/2
```

```
Old MAC: 00e0-ca63-8141
```

```
New MAC: 00e0-ca63-8142
```

```
IPv6 address: 10::2
```

```
System time: 2018-02-02 10:20:30
```

```
Conflict count: 1
```

```
Log suppress count: 0
```

```
Old interface: GigabitEthernet1/0/1
```

```
New interface: GigabitEthernet1/0/2
```

```
Old SVLAN/CVLAN: 100/--
```

```
New SVLAN/CVLAN: 100/--
```

```
Old MAC: 00e0-ca63-8141
```

```
New MAC: 00e0-ca63-8142
```

Table 30 Command output

| Field | Description |
|--------------------|---|
| IPv6 address | IPv6 address of a user. |
| System time | Time when the user IPv6 address conflict occurred. |
| Conflict count | Number of times user IPv6 address conflicts occurred. |
| Log suppress count | Number of times user IPv6 address conflict log generation is suppressed. |
| Old interface | Output interface in the old ND entry. |
| New interface | Output interface in the new ND entry. |
| Old SVLAN/CVLAN | ID of the outer VLAN or inner VLAN in the old ND entry. This field displays hyphens (--) if the ND entry does not belong to any outer VLAN or inner VLAN. |
| New SVLAN/CVLAN | ID of the outer VLAN or inner VLAN in the new ND entry. This field displays hyphens (--) if the ND entry does not belong to any outer VLAN or inner VLAN. |
| Old MAC | MAC address in the old ND entry. |

| | |
|---------|----------------------------------|
| New MAC | MAC address in the new ND entry. |
|---------|----------------------------------|

Related commands

`ipv6 nd user-ip-conflict record enable`

New feature: Enabling recording user port move for ND

Enabling recording user port move for ND

About recording user port move for ND

This feature enables the device to detect and record user port moves. A user port moves if an incoming NA packet has the same source IPv6 address and source MAC address as an existing ND entry but a different port. The device generates a user port move record, logs the move event, and sends the log to the information center. For information about the log destination and output rule configuration in the information center, see the information center in *Network Management and Monitoring Configuration Guide*.

Procedure

1. Enter system view.
`system-view`
2. Enable recording user port move for ND.
`ipv6 nd user-move record enable`

By default, recording user port moves is disabled in ND.

Command reference

ipv6 nd user-move record enable

Use `ipv6 nd user-move record enable` to enable recording user port moves for ND.

Use `undo ipv6 nd user-move record enable` to disable recording user port moves for ND.

Syntax

`ipv6 nd user-move record enable`
`undo ipv6 nd user-move record enable`

Default

Recording user port moves is disabled in ND.

Views

System view

Predefined user roles

network-admin

Usage guidelines

This feature enables the device to detect and record user port move. A user port moves if an incoming NA packet has the same source IPv6 address and source MAC address as an existing ND entry but a different port. The device generates a user port move record, logs the move event, and sends the log to the information center. For information about the log destination and output rule

configuration in the information center, see the information center in *Network Management and Monitoring Configuration Guide*.

An IRF member device can generate a maximum of 10 user port move logs per second. The device suppresses generating user port move logs when their count reaches the upper limit. An IRF member device can save a maximum of 200 user port move records.

When the number of saved user port move records reaches the upper limit, new records overwrite old ones.

Examples

```
# Enable recording user port move for ND.
<Sysname> system-view
[Sysname] ipv6 nd user-move record enable
```

Related commands

display ipv6 nd user-move record

display ipv6 nd user-move record

Use **display ipv6 nd user-move record** to display user port move records.

Syntax

```
display ipv6 nd user-move record [ slot slot-number ]
```

Views

Any view

Predefined user roles

network-admin
network-operator

Parameters

slot *slot-number*: Specifies an IRF member device by its member ID. If you do not specify a member device, this command displays user port move records for all member devices.

Examples

```
# Display all user port move records.
<Sysname> display ipv6 nd user-move record
IPv6 address: 10::1
MAC address: 00e0-ca63-8141
System time: 2018-02-02 11:22:29
Move count: 1
Log suppress count: 0
Before:
  interface: GigabitEthernet1/0/1
  SVLAN/CVLAN: 100/2
After:
  interface: GigabitEthernet1/0/2
  SVLAN/CVLAN: 100/2

IPv6 address: 10::2
MAC address: 00e0-ca63-8142
System time: 2018-02-02 10:20:30
```

```

Move count: 1
Log suppress count: 0
Before:
    interface: GigabitEthernet1/0/1
    SVLAN/CVLAN: 100/--
After:
    interface: GigabitEthernet1/0/2
    SVLAN/CVLAN: 100/--

```

Table 31 Command output

| Field | Description |
|--------------------|---|
| IPv6 address | IPv6 address of the user. |
| MAC address | MAC address of the user. |
| System time | Time when the user port move occurred. |
| Move count | Number of times the user port moved. |
| Log suppress count | Number of times user port move logs are suppressed. |
| Before | Information before the user port move. |
| interface | Interface information in the ND entry. |
| SVLAN/CVLAN | ID of the outer VLAN or inner VLAN in the ND entry. This field displays hyphens (--) if the ND entry does not belong to any outer VLAN or inner VLAN. |
| After | Information after the user port move. |

Related commands

```
ipv6 nd user-move record enable
```

New feature: Enabling ND logging for user online and offline events

Enabling ND logging for user online and offline events

About ND logging for user online and offline events

This feature enables the device to generate user online or offline logs upon such events and send these logs to the information center. For information about the log destination and output rule configuration in the information center, see the information center in *Network Management and Monitoring Configuration Guide*.

Restrictions and guidelines

A higher log output rate consumes more CPU resources. Adjust the log output rate based the CPU performance and usage.

Procedure

1. Enter system view.
system-view
2. Enable ND logging for user online and offline events.
ipv6 nd online-offline-log enable [rate rate]

By default, ND logging for user online and offline events is disabled.

Command reference

ipv6 nd online-offline-log enable

Use **ipv6 nd online-offline-log enable** to enable ND logging for user online and offline events.

Use **undo ipv6 nd online-offline-log enable** to disable ND logging for user online and offline events.

Syntax

```
ipv6 nd online-offline-log enable [ rate rate ]  
undo ipv6 nd online-offline-log enable
```

Default

ND logging for user online and offline events is disabled.

Views

System view

Predefined user roles

network-admin

Parameters

rate *rate*: Specifies the maximum number of logs that can be output per second. The value range is 10 to 300. If you do not specify this option, the maximum log output rate is 100 logs per second.

Usage guidelines

A higher log output rate consumes more CPU resources. Adjust the log output rate based the CPU performance and usage.

Examples

```
# Enable ND logging for user online and offline events, and set the maximum log output rate to 100  
logs per second.  
<Sysname> system-view  
[Sysname] ipv6 nd online-offline-log enable rate 100
```

Related commands

ipv6 neighbor

New feature: Specifying the boot file URL in RA messages

Specifying the boot file URL in RA messages

This release added support for specifying the URL of the boot file in RA messages.

Command reference

ipv6 nd ra boot-file-url

Use `ipv6 nd ra boot-file-url` to specify the URL of the boot file in RA messages.

Use `undo ipv6 nd ra boot-file-url` to restore the default.

Syntax

```
ipv6 nd ra boot-file-url url-string
undo ipv6 nd ra boot-file-url
```

Default

RA messages do not carry the URL of the boot file.

Views

Interface view

Predefined user roles

network-admin

Parameters

url-string: Specifies the URL address of the boot file, a case-sensitive string of 1 to 127 characters. The URL address must be started with `http://`, `https://`, `ftp://`, or `tftp://`.

Usage guidelines

In a data center, a device follows the steps to implement automatic configuration:

1. Obtains an IPv6 address through ND or DHCPv6.
2. Obtains the URL address for downloading the boot file from the DHCPv6 server.
3. Downloads the boot file from the FTP server and installs it.

With the boot file URL specified in RA messages, the device can use the ND protocol to obtain both the IPv6 address and the boot file URL for automatic configuration. DHCPv6 is not required in the network, simplifying network deployment.

Examples

```
# Specify the boot file URL address as tftp://169.254.0.1/file/softimg.iso in RA messages on
VLAN-interface 100.
```

```
<Sysname> system-view
```

```
[Sysname] interface vlan-interface 100
```

```
[Sysname-Vlan-interface100] ipv6 nd ra boot-file-url tftp://169.254.0.1/file/softimg.iso
```

New feature: Setting the aging timer for ND entries in stale state on an interface

Setting the aging timer for ND entries in stale state on an interface

In this release, you can set the aging timer for ND entries in stale state on a per interface basis.

Command reference

ipv6 neighbor timer stale-aging

Use **ipv6 neighbor timer stale-aging** to set the aging timer for ND entries in stale state on an interface.

Use **undo ipv6 neighbor timer stale-aging** to restore the default.

Syntax

```
ipv6 neighbor timer stale-aging aging-time
```

```
undo ipv6 neighbor timer stale-aging
```

Default

The aging timer of ND entries in stale state is not configured on an interface. The aging timer is determined by the configuration of the **ipv6 neighbor stale-aging** command in system view.

Views

Interface view

Predefined user roles

network-admin

Parameters

aging-time: Specifies the aging timer for ND entries in stale state on an interface. The value range is 1 to 1440 minutes.

Usage guidelines

This aging timer applies to ND entries in stale state on the interface. If an ND entry in stale state is not updated before the timer expires, it changes to the delay state. If it is still not updated in 5 seconds, the ND entry changes to the probe state. The device sends an NS message for detection a maximum of three attempts. If no response is received, the device deletes the ND entry.

You can set the aging timer for ND entries in stale state in system view and interface view. For ND entries in stale state on an interface, the aging timer in interface view has higher priority than the aging timer in system view.

Examples

On VLAN-interface 2, set the aging timer to 200 minutes for ND entries in stale state.

```
<Sysname> system-view
```

```
[Sysname] interface vlan-interface 2
```

```
[Sysname-Vlan-interface2] ipv6 neighbor timer stale-aging 200
```

Related commands

```
ipv6 neighbor stale-aging
```

New feature: Specifying DNS server information in RA messages

Specifying DNS server information in RA messages

About specifying DNS server information in RA messages

The DNS server options in RA messages provide DNS server information for IPv6 hosts. A host can obtain both an IPv6 address and the DNS server information through stateless autoconfiguration instead of DHCPv6.

A DNS server option contains information about a DNS server. DNS server options are sorted in ascending order of DNS server sequence number.

After you execute the **ipv6 nd ra dns server** command, the device immediately sends an RA message with existing and new DNS server information. After you execute the **undo ipv6 nd ra dns server** command, the device immediately sends two RA messages. The first RA message carries information about all DNS servers, including the DNS servers that have been removed with their lifetime set to 0 seconds. The second RA message carries information about remaining DNS servers. The device immediately refreshes the RA message advertisement interval for an interface after sending an RA message from the interface.

Restrictions and guidelines

You can configure a maximum of eight DNS servers on an interface.

The default lifetime of the DNS server is three times the maximum interval for advertising RA messages. To set the maximum interval for advertising RA messages, use the **ipv6 nd ra interval** command.

Procedure

1. Enter system view.
system-view
2. Enter interface view.
interface *interface-type interface-number*
3. Specify DNS server information in RA messages.
ipv6 nd ra dns server *ipv6-address* [*seconds* | **infinite**] **sequence** *seqno*

By default, DNS server information is not configured and RA messages do not carry DNS server information.

Command reference

ipv6 nd ra dns server

Use **ipv6 nd ra dns server** to specify DNS server information in RA messages.

Use **undo ipv6 nd ra dns server** to delete DNS server information in RA messages.

Syntax

```
ipv6 nd ra dns server ipv6-address [ seconds | infinite ] sequence seqno  
undo ipv6 nd ra dns server ipv6-address
```

Default

DNS server information is not configured and RA messages do not carry DNS server information.

Views

Interface view

Predefined user roles

network-admin

Parameters

ipv6-address: Specifies the IPv6 address of the DNS server, which must be a global unicast address or a link-local address.

seconds: Specifies the lifetime of the DNS server, in seconds. The value range is 4 to 4294967295. Value 4294967295 indicates that the lifetime of the DNS server is infinite.

infinite: Sets the lifetime of the DNS server to infinite.

sequence seqno: Specifies the sequence number of the DNS server, in the range of 0 to 4294967295. The sequence number for a DNS server must be unique. A smaller sequence number indicates a higher priority.

Usage guidelines

The DNS server option in RA messages provides DNS server information for hosts. A host can obtain both an IPv6 address and the DNS server information through stateless autoconfiguration instead of DHCPv6.

The default lifetime of the DNS server is three times the maximum interval for advertising RA messages. To set the maximum interval for advertising RA messages, use the **ipv6 nd ra interval** command.

You can configure a maximum of eight DNS servers on an interface. A DNS server option contains information about a DNS server. DNS server options are sorted in ascending order of DNS server sequence number.

The sequence number for a DNS server must be unique. A DNS server cannot have multiple sequence numbers. To modify the IPv6 address or sequence number of a DNS server, you must first use the **undo ipv6 nd ra dns server** command to delete information about the DNS server.

After you execute the **ipv6 nd ra dns server** command, the device immediately sends an RA message with existing and new DNS server options. After you execute the **undo ipv6 nd ra dns server** command, the device immediately sends two RA messages. The first RA message carries information about all DNS servers, including the DNS servers that have been removed with their lifetime set to 0 seconds. The second RA message carries information about remaining DNS servers. The device immediately refreshes the RA message advertisement interval for an interface after sending an RA message from the interface.

Examples

Set the DNS server address to 2001:10::100, the lifetime of the DNS server to infinite, and the sequence number to 1 for RA messages advertised on VLAN-interface 100.

```
<Sysname> system-view
[Sysname] interface vlan-interface 100
[Sysname-Vlan-interface100] ipv6 nd ra dns server 2001:10::100 infinite sequence 1
```

Related commands

ipv6 nd ra dns server suppress

ipv6 nd ra interval

New feature: Specifying the DNS domain name suffix information in RA messages

Specifying the DNS domain name suffix information in RA messages

About specifying the DNS domain name suffix information in RA messages

The DNS search list (DNSSL) option in RA messages provides domain name suffix information for IPv6 hosts. A host can obtain both an IPv6 address and the DNS domain name suffix information through stateless autoconfiguration instead of DHCPv6.

A DNSSL option contains information about a DNS domain name suffix. DNSSL options are sorted in ascending order based on the sequence number of each DNS domain name suffix.

After you execute the **ipv6 nd ra dns search-list** command, the device immediately sends an RA message with existing and new DNS domain name suffix information. After you execute the **undo ipv6 nd ra dns search-list** command, the device immediately sends two RA messages. The first RA message carries information about all DNS domain name suffixes, including DNS domain name suffixes that have been removed with their lifetime set to 0 seconds. The second RA message carries information about remaining DNS domain name suffixes. The device immediately refreshes the RA message advertisement interval for an interface after sending an RA message from the interface.

Restrictions and guidelines

You can configure a maximum of eight DNS domain name suffixes on an interface.

The default lifetime of the DNS domain name suffix is three times the maximum interval for advertising RA messages. To set the maximum interval for advertising RA messages, use the **ipv6 nd ra interval** command.

Procedure

1. Enter system view.
system-view
2. Enter interface view.
interface *interface-type* *interface-number*
3. Configure DNS domain name suffix information in RA messages.
ipv6 nd ra dns search-list *domain-name* [*seconds* | **infinite**] **sequence** *seqno*

By default, DNS domain name suffix information is not configured and RA messages do not carry DNS domain name suffix information.

Command reference

ipv6 nd ra dns search-list

Use **ipv6 nd ra dns search-list** to configure DNS domain name suffix information in RA messages.

Use **undo ipv6 nd ra dns search-list** to delete DNS domain name suffix information in RA messages.

Syntax

ipv6 nd ra dns search-list *domain-name* [*seconds* | **infinite**] **sequence** *seqno*

```
undo ipv6 nd ra dns search-list domain-name
```

Default

DNS domain name suffix information is not configured and RA messages do not carry DNS server information.

Views

Interface view

Predefined user roles

network-admin

Parameters

domain-name: Specifies a domain name suffix. It is a dot-separated, case-insensitive string that can include letters, digits, hyphens (-), underscores (_), and dots (.), for example, aabbcc.com. The domain name suffix can include a maximum of 253 characters, and each separated string includes no more than 63 characters.

seconds: Specifies the lifetime of the DNS domain name suffix, in seconds. The value range is 4 to 4294967295. Value 4294967295 indicates that the lifetime of the DNS domain name suffix is infinite.

infinite: Sets the lifetime of the DNS domain name suffix to infinite.

seqno: Specifies the sequence number of the DNS domain name suffix, in the range of 0 to 4294967295. The sequence number for a DNS domain name suffix must be unique. A smaller sequence number indicates a higher priority.

Usage guidelines

The DNS search list (DNSSL) option in RA messages provides DNS domain name suffix information for hosts. A host can obtain both an IPv6 address and the DNS domain name suffix information through stateless autoconfiguration instead of DHCPv6.

The default lifetime of the DNS domain name suffix is three times the maximum interval for advertising RA messages. To set the maximum interval for advertising RA messages, use the **ipv6 nd ra interval** command.

You can configure a maximum of eight DNS domain name suffixes on an interface. Information about a DNS domain name suffix occupies a DNSSL option. All DNSSL options are sorted in ascending order based on the sequence number of each DNS domain name suffix.

The sequence number for a DNS domain name suffix must be unique. A DNS domain name suffix cannot have multiple sequence numbers. To modify a DNS domain name suffix or its sequence number, you must first use the **undo ipv6 nd ra dns search-list** command to delete information about the DNS domain name suffix.

After you execute the **ipv6 nd ra dns search-list** command, the device immediately sends an RA message with existing and new DNS domain name suffix information. After you execute the **undo ipv6 nd ra dns search-list** command, the device immediately sends two RA messages. The first RA message carries information about all DNS domain name suffixes, including DNS domain name suffixes that have been removed with their lifetime set to 0 seconds. The second RA message carries information about remaining DNS domain name suffixes. The device immediately refreshes the RA message advertisement interval for an interface after sending an RA message from the interface.

Examples

Set the DNS domain name suffix to [com](#), the lifetime of the suffix to infinite, and the sequence number to 1 for RA messages advertised on VLAN-interface 100.

```
<Sysname> system-view
```

```
[Sysname] interface vlan-interface 100
```

```
[Sysname-Vlan-interface100] ipv6 nd ra dns search-list com infinite sequence 1
```

Related commands

```
ipv6 nd ra dns search-list suppress
```

```
ipv6 nd ra interval
```

New feature: Suppressing DNS information in RA messages

Suppressing DNS information in RA messages

About suppressing DNS information in RA messages

Perform this task to suppress the device from advertising information about DNS server addresses and domain name suffixes in RA messages.

After you enable suppressing DNS server address advertisement on an interface, the device performs operations as follows:

- If the interface has been configured with DNS server information, the device immediately sends two RA messages. The lifetime for DNS server addresses is 0 seconds in the first message. The second RA message does not carry any DNS server addresses.
- If no DNS server information is configured on the interface, the device does not send RA messages.
- If you configure new DNS server information or removes a DNS server, the device immediately sends an RA message without any DNS server addresses.

After you disable suppression for DNS server addresses on an interface, the device performs operations as follows:

- If the interface has been configured with the DNS server information, the device immediately sends an RA message with the configured DNS server lifetime.
- If no DNS server information is configured on the interface, the device does not send RA messages.

The device immediately refreshes the RA message advertisement interval for an interface after sending an RA message from the interface.

The same suppression mechanism applies when you enable or disable suppression for DNS domain name suffixes in RA messages.

Procedure

1. Enter system view.

```
system-view
```

2. Enter interface view.

```
interface interface-type interface-number
```

3. Enable suppression for DNS server addresses in RA messages.

```
ipv6 nd ra dns server suppress
```

By default, the suppression for DNS addresses in RA messages is disabled.

4. Enable suppression for DNS domain name suffixes in RA messages.

```
ipv6 nd ra dns search-list suppress
```

By default, the suppression for DNS domain name suffixes in RA messages is disabled.

Command reference

ipv6 nd ra dns server suppress

Use **ipv6 nd ra dns server suppress** to enable suppression for DNS server addresses in RA messages.

Use **undo ipv6 nd ra dns server suppress** to disable suppression for DNS server addresses in RA messages.

Syntax

```
ipv6 nd ra dns server suppress
undo ipv6 nd ra dns server suppress
```

Default

Suppression for DNS server addresses in RA messages is disabled.

Views

Interface view

Predefined user roles

network-admin

Usage guidelines

This command enables the device to suppress advertising information about DNS server addresses in RA messages.

RA message advertisement is suppressed by default. To disable RA message advertisement suppression, use the **undo ipv6 nd ra halt** command.

After you enable suppressing DNS server address advertisement on an interface, the device performs operations as follows:

- If the interface has been configured with DNS server information, the device immediately sends two RA messages. The lifetime for DNS server addresses is 0 seconds in the first message. The second RA message does not carry the DNS server option.
- If no DNS server information is configured on the interface, the device does not send RA messages.
- If you configure new DNS server information or removes a DNS server, the device immediately sends an RA message without any DNS server addresses.

After you disable suppressing DNS server address advertisement on an interface, the device performs operations as follows:

- If the interface has been configured with the DNS server information, the device immediately sends an RA message with the configured DNS server lifetime.
- If no DNS server information is configured on the interface, the device does not send RA messages.

The device immediately refreshes the RA message advertisement interval for an interface after sending an RA message from the interface.

Examples

Enable suppression for DNS server addresses in RA messages on VLAN-interface 100.

```
<Sysname> system-view
[Sysname] interface vlan-interface 100
[Sysname-Vlan-interface100] ipv6 nd ra dns server suppress
```

Related commands

`ipv6 nd ra dns server`

ipv6 nd ra dns search-list suppress

Use `ipv6 nd ra dns search-list suppress` to enable suppression for DNS domain name suffixes in RA messages.

Use `undo ipv6 nd ra dns search-list suppress` to disable suppression for DNS domain name suffixes in RA messages.

Syntax

`ipv6 nd ra dns search-list suppress`

`undo ipv6 nd ra dns search-list suppress`

Default

Suppression for DNS domain name suffixes in RA messages is disabled.

Views

Interface view

Predefined user roles

network-admin

Usage guidelines

This command enables the device to suppress advertising information about domain name suffixes in RA messages.

RA message advertisement is suppressed by default. To disable RA message advertisement suppression, use the `undo ipv6 nd ra halt` command.

After you enable suppressing DNS domain name suffix advertisement on an interface, the device performs operations as follows:

- If the interface has been configured with DNS domain name suffix information, the device immediately sends two RA messages. The lifetime for DNS domain name suffixes is 0 seconds in the first message. The second RA message does not carry any DNS domain name suffixes.
- If no DNS domain name suffix information is configured on the interface, the device does not send RA messages.
- If you configure new DNS domain name suffix information or removes a DNS domain name suffix after the suppression is enabled, the device immediately sends an RA message without any DNS domain name suffixes.

After you disable suppressing DNS domain name suffix advertisement on an interface, the device performs operations as follows:

- If the interface has been configured with the DNS domain name suffix information, the device immediately sends an RA message with the configured DNS domain name suffix lifetime.
- If no DNS domain name suffix information is configured on the interface, the device does not send RA messages.

The device immediately refreshes the RA message advertisement interval for an interface after sending an RA message from the interface.

Examples

Enable suppression for DNS domain name suffixes in RA messages on VLAN-interface 100.

```
<Sysname> system-view
```

```
[Sysname] interface vlan-interface 100
```

```
[Sysname-Vlan-interface100] ipv6 nd ra dns search-list suppress
```

Related commands

```
ipv6 nd ra dns search-list
```

New feature: Enabling the DHCPv6 relay agent to advertise host routes for assigned IPv6 addresses

Enabling the DHCPv6 relay agent to advertise host routes for assigned IPv6 addresses

About enabling the DHCPv6 relay agent to advertise host routes for assigned IPv6 addresses

In a network where ND cannot resolve global unicast addresses, network devices cannot generate ND entries for all global unicast addresses. If a DHCPv6 client obtains a global unicast address, the neighboring devices do not have the ND entries for this global unicast address, thus cannot forward the packets destined for the client. To resolve this problem, enable the DHCPv6 relay agent to advertise host routes for assigned IPv6 addresses in DHCP replies. The advertised route information is as follows:

- The destination IP address is the assigned IPv6 address.
- The next hop is the link-local address of the DHCPv6 client.
- The output interface is the interface that forwards the reply.

After the relay agent receives a packet destined for the assigned IPv6 address, the relay agent looks up the routing table for the next hop. ND resolution can succeed because the next hop is the link-local address of the client. The relay agent searches the ND table for the MAC address of the client based on the next hop and then forwards the packet.

Restrictions and guidelines

Before using this feature on the DHCPv6 relay agent, enable the DHCPv6 relay agent to record DHCPv6 relay entries first.

Procedure

1. Enter system view.
system-view
2. Enable the DHCP relay agent to advertise host routes for assigned IPv6 addresses.
ipv6 dhcp advertise address-route
By default, the DHCPv6 relay agent does not advertise host routes for assigned IPv6 addresses.

Command reference

Use **ipv6 dhcp advertise address-route** to enable the DHCPv6 relay agent to advertise host routes for IPv6 addresses assigned to DHCP clients.

Use **undo ipv6 dhcp advertise address-route** to disable the DHCPv6 relay agent from advertising host routes for assigned IPv6 addresses.

Syntax

```
ipv6 dhcp advertise address-route
```

```
undo ipv6 dhcp advertise address-route
```

Default

The DHCPv6 relay agent does not advertise host routes for assigned IPv6 addresses.

Views

System view

Predefined user roles

network-admin

Usage guidelines

In a network where ND cannot resolve global unicast addresses, network devices cannot generate ND entries for all global unicast addresses. If a DHCPv6 client obtains a global unicast address, the neighboring devices do not have the ND entries for this global unicast address, thus cannot forward the packets destined for the client. To resolve this problem, enable the DHCPv6 relay agent to advertise host routes for assigned IPv6 addresses in DHCP replies. The advertised route information is as follows:

- The destination IP address is the assigned IPv6 address.
- The next hop is the link-local address of the DHCPv6 client.
- The output interface is the interface that forwards the reply.

After the relay agent receives a packet destined for the assigned IPv6 address, the relay agent looks up the routing table for the next hop. ND resolution can succeed because the next hop is the link-local address of the client. The relay agent searches the ND table for the MAC address of the client based on the next hop and then forwards the packet.

Before using this command on the DHCPv6 relay agent, enable the DHCPv6 relay agent to record DHCPv6 relay entries first.

Examples

```
# Enable the DHCPv6 relay agent to advertise host routes for assigned IPv6 addresses.
```

```
<Sysname> system-view
```

```
[Sysname] ipv6 dhcp advertise address-route
```

New feature: Enabling IPv6 address binding conversion for IP source guard

Enabling IPv6 address binding conversion for IP source guard

About enabling IPv6 address binding conversion for IP source guard

In a VXLAN network, the VXLAN IP gateway acts as the DHCPv6 server to assign IPv6 addresses to users and reports user information to the controller. Only IP source guard entries are supported in user information report. This feature enables the DHCP server to generate dynamic IP source guard bindings based on IPv6 address bindings of users. The controller manages and monitors these users based on the dynamic IP source guard bindings.

For more information about IP source guard, see *Security Configuration Guide*.

Restrictions and guidelines

If the device does not have enough storage space, execute the **undo ipv6 dhcp server entry-convert enable** command to disable this feature. The system does not delete IP source guard bindings that have been converted before you disable this feature.

Procedure

1. Enter system view.
system-view
2. Enable IPv6 address binding conversion for IP source guard.
ipv6 dhcp server entry-convert enable
By default, IPv6 address binding conversion for IP source guard is disabled.

Command reference

ipv6 dhcp server entry-convert enable

Use **ipv6 dhcp server entry-convert enable** to enable IPv6 address binding conversion for IP source guard.

Use **undo ipv6 dhcp server entry-convert enable** to disable IPv6 address binding conversion for IP source guard.

Syntax

```
ipv6 dhcp server entry-convert enable
undo ipv6 dhcp server entry-convert enable
```

Default

IPv6 address binding conversion for IP source guard is disabled.

Views

System view

Predefined user roles

network-admin

Usage guidelines

In a VXLAN network, the VXLAN IP gateway acts as the DHCPv6 server to assign IPv6 addresses to users and reports user information to the controller. Only IP source guard entries are supported in user information report. This feature enables the DHCP server to generate dynamic IP source guard bindings based on IPv6 address bindings of users. The controller manages and monitors these users based on the dynamic IP source guard bindings.

If the device does not have enough storage space, execute the **undo ipv6 dhcp server entry-convert enable** command to disable this feature. The system does not delete IP source guard bindings that have been converted before you disable this feature.

For more information about IP source guard, see *Security Configuration Guide*.

Examples

Enable IPv6 address binding conversion for IP source guard.

```
<Sysname> system-view
[Sysname] ipv6 dhcp server entry-convert enable
```

New feature: Configuring DHCPv6 guard

About DHCPv6 guard

The DHCPv6 guard feature filters DHCPv6 Advertise and Reply messages by using DHCPv6 guard policies to make sure DHCPv6 clients obtain addresses/prefixes from authorized DHCPv6 servers. To provide finer level of filtering granularity, you can specify the following parameters for a DHCPv6 guard policy:

- Device role of the device that attached to the target interface or VLAN. The interface or VLAN to which the DHCPv6 guard policy is applied is called the target interface or VLAN.
- DHCPv6 server match criterion.
- Match criterion for IPv6 addresses/prefixes assigned by DHCPv6 servers.
- Allowed DHCPv6 server preference range.

To meet requirements of DHCPv6 clients in different locations, apply DHCPv6 guard policies to different interfaces or VLANs on the same device.

DHCPv6 guard operating mechanism

Upon receiving a DHCPv6 Solicit or Request message, the DHCPv6 guard device forwards the message without performing the DHCPv6 guard policy check.

When receiving a DHCPv6 reply, the DHCPv6 guard device performs the DHCPv6 guard policy check in the following order:

1. Examines whether the receiving port is a trusted port. The device forwards the message if the message is from the a trusted port.

Configure trusted ports in a DHCPv6 guard policy only in one of the following conditions:

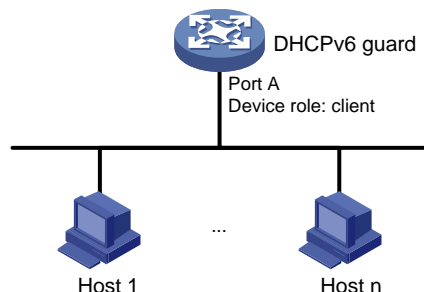
- The port to which the DHCPv6 guard policy applies is connected to an authorized server.
- All ports in the VLAN to which the DHCPv6 guard policy applies are connected to authorized servers.

2. Examines the message based on the device role:

- If the message is received from the device with the DHCPv6 client device role, the device drops the message.

If the interface to which the DHCPv6 guard policy applies is not connected to any authorized DHCPv6 servers, set the device role to **client** for the policy, as shown in [Figure 3](#).

Figure 3 Setting the device role to client

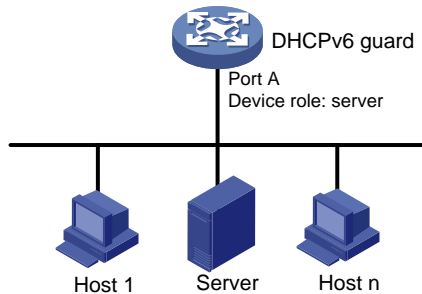


- If the message is received from the device with the DHCPv6 server device role, the device examines the message as follows:
 - For an Advertise message, the message passes the policy check if the source IP address in the message is permitted by the ACL and the server preference is in the match range.

- For a Reply message, the message passes the policy check if the assigned IPv6 addresses/prefixes in the message are permitted by the ACL.

If the interface to which the DHCPv6 guard policy applies is connected to an authorized DHCPv6 server, set the device role to **server** for the policy, as shown in [Figure 4](#).

Figure 4 Setting the device role to server



The device forwards the reply after the message passes the DHCPv6 guard policy check.

Restrictions and guidelines: DHCPv6 guard configuration

The DHCPv6 guard feature operates correctly only when the device is located between the DHCPv6 client and the DHCPv6 server or between the DHCPv6 client and the DHCPv6 relay agent. If the device is located between the DHCPv6 server and the DHCPv6 relay agent, the DHCPv6 guard feature cannot operate correctly.

When the DHCPv6 guard feature is configured on a DHCPv6 snooping device, both features can take effect. The device forwards DHCPv6 reply packets received on a DHCPv6 snooping trusted port only if they pass the DHCPv6 guard check. These packets are dropped if they fail the DHCPv6 guard check.

DHCPv6 guard tasks at a glance

To configure DHCPv6 guard, perform the following tasks:

1. [Configuring a DHCPv6 guard policy](#)
2. Applying the DHCPv6 guard policy

Choose the following tasks as needed:

- [Applying a DHCPv6 guard policy to an interface](#)
- [Applying a DHCPv6 guard policy to a VLAN](#)

If DHCPv6 guard policies are applied to both an interface and the VLAN of the interface, the interface-specific policy is used on the interface.

Configuring a DHCPv6 guard policy

1. Enter system view.
system-view
2. Create a DHCPv6 guard policy and enter its view.
ipv6 dhcp guard policy *policy-name*
3. Specify the role of the device attached to the target interface or VLAN.
device-role { **client** | **server** }

By default, the device role is DHCPv6 client for the device attached to the target interface or VLAN.

4. Configure a DHCPv6 guard policy.
 - Configure a DHCPv6 server match criterion.
`if-match server acl { acl-number | name acl-name }`
 By default, no DHCPv6 server match criterion is configured, and all DHCPv6 servers are authorized.
 - Configure a match criterion for the assigned IPv6 addresses/prefixes.
`if-match reply acl { acl-number | name acl-name }`
 By default, no match criterion is configured for the assigned IPv6 addresses/prefixes, and all assigned IPv6 addresses/prefixes can pass the address/prefix check.
 - Configure an allowed DHCPv6 server preference range.
`preference { max max-value | min min-value } *`
 By default, no DHCPv6 server preference range is configured, and DHCPv6 servers with preferences 1 to 255 can pass the preference check.
 - Configure the port to which the policy applies as a trusted port for the policy.
`trust port`
 By default, no trusted port is configured for a DHCPv6 guard policy.

Applying a DHCPv6 guard policy to an interface

1. Enter system view.
`system-view`
2. Enter Layer 2 interface view.
`interface interface-type interface-number`
3. Apply a DHCPv6 guard policy to the interface.
`ipv6 dhcp guard apply policy policy-name`
 By default, no DHCPv6 guard policy is applied to the interface.

Applying a DHCPv6 guard policy to a VLAN

1. Enter system view.
`system-view`
2. Create a VLAN and enter its view.
`vlan vlan-number`
3. Apply a DHCPv6 guard policy to the VLAN.
`ipv6 dhcp guard apply policy policy-name`
 By default, no DHCPv6 guard policy is applied to the VLAN.

Display and maintenance commands for DHCPv6 guard

Execute **display** commands in any view.

| Task | Command |
|--|---|
| Display information about DHCPv6 guard policies. | <code>display ipv6 dhcp guard policy</code> [<i>policy-name</i>] |

DHCPv6 guard configuration examples

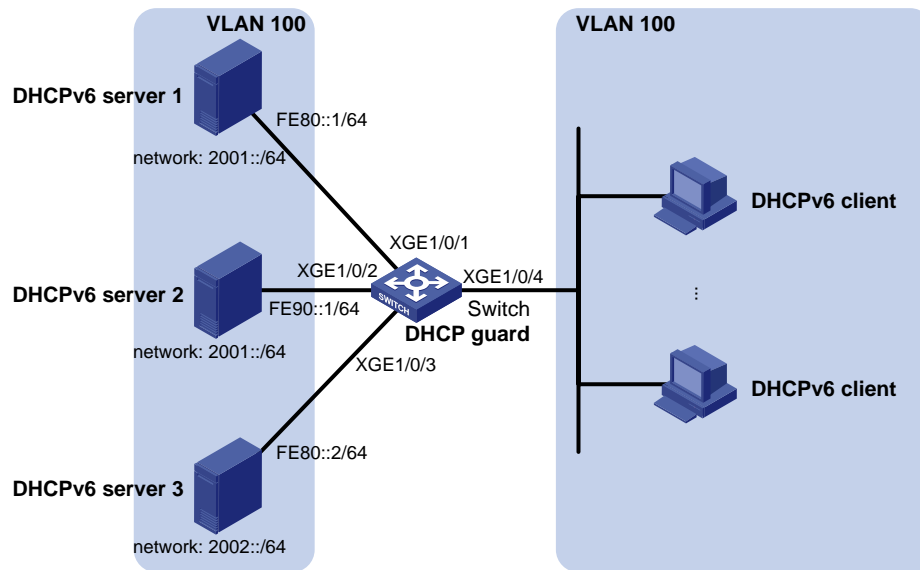
Example: Configuring DHCPv6 guard

Network configuration

As shown in Figure 5, all DHCPv6 servers and clients are in VLAN 100. The assignable IPv6 address ranges on the DHCPv6 server 1, server 2, and server 3 are 2001::/64, 2001::/64, and 2002::/64, respectively.

Configure DHCPv6 guard on the switch, so that the switch forwards only DHCPv6 replies with the source IPv6 address in the range of FE80::/12 and assigned prefixes in the range of 2001::/16.

Figure 5 Network diagram



Procedure

Before you configure DHCPv6 guard, complete the configuration on DHCPv6 servers.

Create VLAN 100, and assign Ten-GigabitEthernet1/0/1, Ten-GigabitEthernet1/0/2, Ten-GigabitEthernet1/0/3, and Ten-GigabitEthernet1/0/4 to VLAN 100.

```
<Switch> system-view
[Switch] vlan 100
[Switch-vlan100] port ten-gigabitethernet 1/0/1 to ten-gigabitethernet 1/0/4
[Switch-vlan100] quit
```

Create an IPv6 basic ACL numbered 2001.

```
[Switch] acl ipv6 number 2001
```

Create rule 1 to permit only packets with source IPv6 addresses in the range of FE80::/12.

```
[Switch-acl-ipv6-basic-2001] rule 1 permit source fe80:: 12
[Switch-acl-ipv6-basic-2001] quit
```

Create an IPv6 basic ACL numbered 2002.

```
[Switch] acl ipv6 number 2002
```

Create rule 1 to permit only packets with source IPv6 addresses in the range of 2001::/16.

```
[Switch-acl-ipv6-basic-2002] rule 1 permit source 2001:: 16
[Switch-acl-ipv6-basic-2002] quit
```

Create DHCPv6 guard policy named p1.

```

[Switch] ipv6 dhcp guard policy p1

# Set the device role to the DHCPv6 server for the device attached to the target VLAN.
[Switch-dhcp6-guard-policy-p1] device-role server

# Specify ACL 2001 to match DHCPv6 servers.
[Switch-dhcp6-guard-policy-p1] if-match server acl 2001

# Specify ACL 2002 to match IPv6 addresses/prefixes assigned by DHCPv6 servers.
[Switch-dhcp6-guard-policy-p1] if-match reply acl 2002
[Switch-dhcp6-guard-policy-p1] quit

# Create DHCPv6 guard policy named p2.
[Switch] ipv6 dhcp guard policy p2

# Set the device role to the DHCPv6 client for the device attached to the target interface.
[Switch-dhcp6-guard-policy-p2] device-role client
[Switch-dhcp6-guard-policy-p2] quit

# Apply DHCPv6 guard policy p1 to VLAN 100.
[Switch] vlan 100
[Switch-vlan100] ipv6 dhcp guard apply policy p1
[Switch-vlan100] quit

# Apply DHCPv6 guard policy p2 to Ten-GigabitEthernet1/0/4.
[Switch]interface ten-gigabitethernet 1/0/4
[Switch-Ten-GigabitEthernet1/0/4] ipv6 dhcp guard apply policy p2
[Switch-Ten-GigabitEthernet1/0/4] quit

```

Verifying the configuration

Verify that the switch forwards DHCPv6 replies with the source IPv6 address in the range of FE80::/12 and the assigned IPv6 prefixes in the range of 2001::/16. The switch forwards DHCPv6 replies from the DHCPv6 server 1 and drops replies from DHCPv6 server 2 and server 3.

DHCPv6 guard commands

The DHCPv6 guard feature operates correctly only when the device is located between the DHCPv6 client and the DHCPv6 server or between the DHCPv6 client and the DHCPv6 relay agent. If the device is located between the DHCPv6 server and the DHCPv6 relay agent, the DHCPv6 guard feature cannot operate correctly.

When the DHCPv6 guard feature is configured on a DHCPv6 snooping device, both features can take effect. The device forwards DHCPv6 reply packets received on a DHCP snooping trusted port only if they pass the DHCPv6 guard check. These packets are dropped if they fail the DHCPv6 guard check.

device-role

Use **device-role** to set the role of the device attached to the target interface or VLAN.

Use **undo device-role** to restore the default.

Syntax

```

device-role { client | server }
undo device-role

```

Default

The role is DHCPv6 client for the device attached to the target interface or VLAN.

Views

DHCPv6 guard policy view

Predefined user roles

network-admin

Parameters

client: Sets the device role to DHCPv6 client.

server: Sets the device role to DHCPv6 server.

Usage guidelines

The target interface or VLAN refers to the interface or VLAN to which a DHCPv6 guard policy is applied. The device makes forwarding decisions based on the device role as follows:

- Drops DHCPv6 replies received from the device with the device role of DHCPv6 client.
- Forwards DHCP replies received from the device with the device role of DHCPv6 server only if the packets pass the DHCPv6 guard check.

If the target interface or VLAN is attached to an authorized DHCPv6 server, set the device role to DHCPv6 server for the authorized DHCPv6 server. If no authorized DHCP servers are attached to the target interface or VLAN, set the device role to DHCPv6 client for devices attached to the target interface or VLAN.

The **trust port** command has a higher priority than the **device-role** command. If you configure both commands for a DHCPv6 guard policy, the **trust port** command takes effect.

Examples

Set the role to DHCPv6 server for the device attached to the target interface or VLAN.

```
<Sysname> system-view
[Sysname] ipv6 dhcp guard policy p1
[Sysname-dhcp6-guard-policy-p1] device-role server
```

display ipv6 dhcp guard policy

Use **display ipv6 dhcp guard policy** to display information about DHCPv6 guard policies.

Syntax

display ipv6 dhcp guard policy [*policy-name*]

Views

Any view

Predefined user roles

network-admin

network-operator

Parameters

policy-name: Displays detailed information about a DHCPv6 guard policy. This argument specifies the name of a DHCPv6 guard policy, a case-insensitive string of 1 to 63 characters. If you do not specify this argument, the command displays brief information about all DHCPv6 guard policies.

Examples

Display detailed information about DHCPv6 guard policy **p1**.

```
<Sysname> display ipv6 dhcp guard policy p1
Guard policy: p1
```

```

Device-role: Server
Trusted port: No
Server preference min value: 23
Server preference max value: 45
Server rule: ACL sed
Reply rule: ACL 3434
Applied to interfaces: XGE1/0/1, XGE1/0/2
Applied to VLANs: 100

```

Display brief information about all DHCPv6 guard policies.

```
<Sysname> display ipv6 dhcp guard policy
```

```
Guard policy: p1
```

```

Device-role: server
Trusted port: No
Server preference min value: 23
Server preference max value: 45
Server rule: ACL sed
Reply rule: ACL 3434

```

```
Guard policy: p2
```

```

Device-role: Server
Trusted port: Yes
Server preference min value: 12
Server preference max value: 34

```

Table 32 Command output

| Field | Description |
|-----------------------------|---|
| Guard policy | DHCPv6 guard policy name. |
| Device-role | Device role: <ul style="list-style-type: none"> Client—DHCPv6 client role. Server—DHCPv6 server role. |
| Trusted port | Whether the trusted port is configured for the guard policy. |
| Server preference min value | Minimum preference value of the DHCPv6 server. This field is displayed only when the preference min command is configured. |
| Server preference max value | Maximum preference value of the DHCPv6 server. This field is displayed only when the preference max command is configured. |
| Server rule | DHCP server address match criterion. This field is displayed only when the if-match server acl command is configured. |
| Reply rule | Match criterion for the assigned IPv6 addresses/prefixes. This field is displayed only when the if-match reply acl command is configured. |
| Applied to interfaces | Interfaces to which the DHCPv6 guard policy is applied. Interfaces are separated by commas (.). This field is not displayed when the command displays brief information about DHCPv6 guard policies. |
| Applied to VLANs | VLANs to which the DHCPv6 guard policy is applied. VLANs are separated by commas (.). This field is not displayed when the command displays brief information about DHCPv6 guard policies. |

Related commands

`ipv6 dhcp guard policy`

if-match reply acl

Use `if-match reply acl` to configure a match criterion for IPv6 addresses/prefixes assigned by a DHCPv6 server.

Use `undo if-match server acl` to restore the default.

Syntax

```
if-match reply acl { acl-number | name acl-name }  
undo if-match reply acl
```

Default

No match criterion is configured for the assigned IPv6 addresses/prefixes, and all assigned IPv6 addresses/prefixes can pass the address/prefix check.

Views

DHCPv6 guard policy view

Predefined user roles

network-admin

Parameters

acl-number: Specifies an ACL number. The value range for this argument is as follows:

- 2000 to 2999 for a basic ACL.
- 3000 to 3999 for an advanced ACL.

name *acl-name*: Specifies a basic or advanced ACL by its name, a case-insensitive string of 1 to 63 characters. The ACL name must start with an English letter and to avoid confusion, it cannot be **all**.

Usage guidelines

The device uses the source IPv6 address attributes in the specified ACL to match the assigned IPv6 address/prefix in the received DHCPv6 Reply message.

- If the assigned IPv6 address/prefix matches a permit statement in the ACL, the device forwards the Reply message. If the assigned IPv6 address/prefix does not match the ACL, the device drops the Reply message.
- If the ACL does not have any source IPv6 address attributes, all DHCPv6 Reply messages fail the address/prefix check and are dropped.
- If the ACL does not exist or does not have any rules, all DHCPv6 Reply messages can pass the check.

If you execute this command multiple times for a DHCPv6 guard policy, the most recent configuration takes effect.

Examples

Specify ACL 2233 to match IPv6 addresses/prefixes assigned by a DHCPv6 server.

```
<Sysname> system-view
```

```
[Sysname] ipv6 dhcp guard policy p1
```

```
[Sysname-dhcp6-guard-policy-p1] if-match reply acl 2233
```

Related commands

`acl` (*ACL and QoS Command Reference*)

rule (IPv6 advanced ACL view) (*ACL and QoS Command Reference*)

rule (IPv6 basic ACL view) (*ACL and QoS Command Reference*)

if-match server acl

Use **if-match server acl** to configure a DHCPv6 server match criterion

Use **undo if-match server acl** to restore the default.

Syntax

```
if-match server acl { acl-number | name acl-name }  
undo if-match server acl
```

Default

No DHCP server match criterion is configured, and all DHCPv6 servers are authorized.

Views

DHCPv6 guard policy view

Predefined user roles

network-admin

Parameters

acl-number: Specifies an ACL number. The value range for this argument is as follows:

- 2000 to 2999 for a basic ACL.
- 3000 to 3999 for an advanced ACL.

name *acl-name*: Specifies a basic or advanced ACL by its name, a case-insensitive string of 1 to 63 characters. The ACL name must start with an English letter and to avoid confusion, it cannot be **all**.

Usage guidelines

The device uses the source IPv6 address attributes in the specified ACL to match the source IPv6 address in the received DHCPv6 Advertise message.

- If the source IPv6 address matches a permit statement in the ACL, the device continues to use other criterion to verify the message. If the source IPv6 address does not match the ACL, the device drops the Advertise message.
- If the ACL does not have any source IPv6 address attributes, all DHCPv6 Advertise messages fail the address check and are dropped.
- If the ACL does not exist or does not have any rules, all DHCPv6 Advertise messages can pass the check.

If you execute this command multiple times for a DHCPv6 guard policy, the most recent configuration takes effect.

Examples

Specify ACL 2323 to match DHCPv6 servers.

```
<Sysname> system-view  
[Sysname] ipv6 dhcp guard policy p1  
[Sysname-dhcp6-guard-policy-p1] if-match server acl 2323
```

Related commands

acl (*ACL and QoS Command Reference*)

rule (IPv6 advanced ACL view) (*ACL and QoS Command Reference*)

rule (IPv6 basic ACL view) (*ACL and QoS Command Reference*)

ipv6 dhcp guard apply policy

Use **ipv6 dhcp guard apply policy** to apply a DHCPv6 guard policy to an interface or a VLAN.

Use **undo ipv6 dhcp guard apply policy** to restore the default.

Syntax

```
ipv6 dhcp guard apply policy policy-name  
undo ipv6 dhcp guard apply policy
```

Default

No DHCPv6 guard policy is applied to an interface or VLAN.

Views

Interface view

VLAN view

Predefined user roles

network-admin

Parameters

policy-name: Specifies a DHCPv6 guard policy name, a case-insensitive string of 1 to 63 characters.

Usage guidelines

The DHCPv6 guard policy applied to an interface checks all incoming DHCP replies if the interface is not configured as a trusted port for the DHCPv6 guard policy.

The DHCPv6 guard policy applied to a VLAN checks all incoming DHCP replies if the interfaces in the VLAN are not configured as trusted ports for the DHCPv6 guard policy.

If you apply a nonexistent DHCPv6 guard policy to an interface or VLAN, the device forwards received DHCPv6 replies without check.

If you execute this command multiple times, the most recent configuration takes effect.

Examples

Apply DHCPv6 guard policy **p1** to Ten-GigabitEthernet1/0/1.

```
<Sysname> system-view
```

```
[Sysname] interface ten-gigabitethernet 1/0/1
```

```
[Sysname-Ten-GigabitEthernet1/0/1] ipv6 dhcp guard apply policy p1
```

Apply DHCPv6 guard policy **p1** to VLAN 100.

```
<Sysname> system-view
```

```
[Sysname] vlan 100
```

```
[Sysname-vlan100] ipv6 dhcp guard apply policy p1
```

Related commands

```
ipv6 dhcp guard policy
```

ipv6 dhcp guard policy

Use **ipv6 dhcp guard policy** to create a DHCPv6 guard policy and enter its view, or enter the view of an existing DHCPv6 guard policy.

Use **undo ipv6 dhcp guard policy** to delete a DHCPv6 guard policy.

Syntax

```
ipv6 dhcp guard policy policy-name  
undo ipv6 dhcp guard policy policy-name
```

Default

No DHCPv6 guard policies exist.

Views

System view

Predefined user roles

network-admin

Parameters

policy-name: Specifies a DHCPv6 guard policy name, a case-insensitive string of 1 to 63 characters.

Usage guidelines

To provide finer level of filtering granularity, you can specify the following parameters for a DHCPv6 guard policy:

- Device role of the device that attached to the target interface or VLAN.
- DHCPv6 server match criterion.
- Match criterion for IPv6 addresses/prefixes assigned by DHCPv6 servers.
- Allowed DHCPv6 server preference range.

The DHCPv6 guard feature runs correctly after you create a DHCPv6 guard policy and apply it to a VLAN or an interface. The DHCPv6 guard feature determines whether to forward DHCP replies based on the match criteria. Only packets that match all criteria are forwarded.

Examples

Create DHCPv6 guard policy **p1** and enter its view.

```
<Sysname> system-view  
[Sysname] ipv6 dhcp guard policy p1  
[Sysname-dhcp6-guard-policy-p1]
```

Related commands

```
display ipv6 dhcp guard policy  
ipv6 dhcp guard apply policy
```

preference

Use **preference** to specify an allowed DHCPv6 server preference range.

Use **undo preference** to restore the maximum or minimum preference to the default value.

Syntax

```
preference { max max-value | min min-value } *  
undo preference [ max | min ]
```

Default

No DHCPv6 server preference range is configured, and DHCPv6 servers with preferences 1 to 255 can pass the preference check.

Views

DHCPv6 guard policy view

Predefined user roles

network-admin

Parameters

max *max-value*: Specifies the maximum value of the DHCPv6 server preference, in the range of 1 to 255. The default is 255.

min *min-value*: Specifies the minimum value of the DHCPv6 server preference, in the range of 1 to 255. The default is 1. The minimum value cannot be higher than the maximum value.

Usage guidelines

The device uses the specified range to match the DHCPv6 server preference in the received DHCPv6 Advertise message.

- If the DHCPv6 server preference is in the allowed range, the device continues to use other criterion to further match the message.
- If the DHCPv6 server preference in the Advertise message is beyond the allowed range or the message does not carry the preference, the device drops the message.

When a device acts as a DHCPv6 server, use the **ipv6 dhcp server preference** command to set the preference of the DHCPv6 server.

If you execute this command multiple times for a DHCPv6 guard policy, the most recent configuration takes effect.

Examples

Set the allowed range to 1 to 100 for the DHCPv6 server preference.

```
<Sysname> system-view
```

```
[Sysname] ipv6 dhcp guard policy p1
```

```
[Sysname-dhcp6-guard-policy-p1] preference max 100 min 1
```

Related commands

ipv6 dhcp server

trust port

Use **trust port** to configure the port to which the DHCPv6 guard policy applies as a trusted port for the policy.

Use **undo trust port** to restore the default.

Syntax

trust port

undo trust port

Default

No trusted port is configured for a DHCPv6 guard policy.

Views

DHCPv6 guard policy view

Predefined user roles

network-admin

Usage guidelines

After you configure this command for a DHCPv6 guard policy, the interface and all interfaces in the VLAN to which the DHCPv6 guard policy is applied are trusted ports. The device forwards received DHCP replies on the trusted ports without check.

The **trust port** command has a higher priority than the **device-role** command. If you configure both commands for a DHCPv6 guard policy, the **trust port** command takes effect.

Examples

Configure the port as a trusted port for the DHCPv6 guard policy.

```
<Sysname> system-view
[Sysname] ipv6 dhcp guard policy p1
[Sysname-dhcp6-guard-policy-p1] trust port
```

New feature: Ignoring router IDs during optimal route selection

Ignoring router IDs during optimal route selection

About ignoring router IDs during optimal route selection

By default, BGP compares router IDs during optimal route selection. If multiple routes to the same destination are available, BGP selects the route with the smallest router ID as the optimal route.

Perform this task to enable BGP to ignore router IDs during optimal route selection.

Procedure

1. Enter system view.
system-view
2. Enter BGP instance view or BGP-VPN instance view.
 - o Enter BGP instance view.
bgp *as-number* [**instance** *instance-name*]
 - o Enter BGP-VPN instance view.
bgp *as-number* [**instance** *instance-name*]
ip vpn-instance *vpn-instance-name*
3. Configure BGP to ignore router IDs during optimal route selection.
bestroute router-id-ignore
By default, BGP compares router IDs during optimal route selection.

Command reference

bestroute router-id-ignore

Use **bestroute router-id-ignore** to configure BGP to ignore router IDs during optimal route selection.

Use **undo bestroute router-id-ignore** to restore the default.

Syntax

bestroute router-id-ignore

```
undo bestroute router-id-ignore
```

Default

By default, BGP compares router IDs during optimal route selection. If multiple routes to the same destination are available, BGP selects the route with the smallest router ID as the optimal route.

Views

BGP instance view

BGP-VPN instance view

Predefined user roles

network-admin

Examples

In BGP instance view, configure BGP to ignore router IDs during optimal route selection.

```
<Sysname> system-view
```

```
[Sysname] bgp 1
```

```
[Sysname-bgp-default] bestroute router-id-ignore
```

New feature: Redistributing the local network specified in the public instance or a VPN instance

Redistributing the local network specified in the public instance or a VPN instance

About redistributing the local network specified in the public instance or a VPN instance

Perform this task to redistribute the local network specified in the public instance or a VPN instance to the BGP routing table, so BGP can advertise the network. The ORIGIN attribute is IGP for routes advertised in this way.

The local network must exist and be active in the routing table of the public instance or VPN instance.

Procedure

1. Enter system view.
system-view
2. Enter public instance IPv4 VPN view or VPN instance IPv4 VPN view.
 - o Enter public instance IPv4 VPN view.
ip public-instance
address-family ipv4
 - o Enter VPN instance IPv4 VPN view.
ip vpn-instance *vpn-instance-name*
address-family ipv4
3. Specify a local network in the public instance or VPN instance.
network *ipv4-address* [*mask-length* | *mask*]
By default, no local network is specified in the public instance or VPN instance.
For more information about this command, see MPLS L3VPN in *MPLS Command Reference*.
4. Execute the quit command twice to return to system view.
quit

5. Enter BGP IPv4 unicast address family view or BGP-VPN IPv4 unicast address family view.
 - Enter BGP IPv4 unicast address family view.
`bgp as-number [instance instance-name]`
`address-family ipv4 [unicast]`
 - Enter BGP-VPN IPv4 unicast address family view.
`bgp as-number [instance instance-name]`
`ip vpn-instance vpn-instance-name`
`address-family ipv4 [unicast]`
6. Redistribute the local network specified in the public instance or VPN instance.
`import-route local-aggregate [route-policy route-policy-name]`
 By default, BGP does not redistribute the local network specified in the public instance or any VPN instances.

Command reference

import-route

Use **import-route local-aggregate** to redistribute the local network specified in the public instance or a VPN instance.

Use **undo import-route local-aggregate** to remove the configuration.

Syntax

import-route local-aggregate [*route-policy route-policy-name*]

undo import-route local-aggregate

Default

BGP does not redistribute the local network specified in the public instance or any VPN instances.

Views

BGP IPv4 unicast address family view

BGP-VPN IPv4 unicast address family view

Predefined user roles

network-admin

Usage guidelines

The ORIGIN attribute is IGP for routes redistributed by the **import-route local-aggregate** command.

Examples

In BGP IPv4 unicast address family view, configure BGP to redistribute local network 10.0.0.0 from VPN instance **vpn1**.

```
<sysname> system-view
[sysname] ip vpn-instance vpn1
[sysname-vpn-instance-vpn1] address-family ipv4
[sysname-vpn-ipv4-vpn1] network 10.0.0.0 255.255.0.0
[sysname-vpn-ipv4-vpn1] quit
[sysname-vpn-instance-vpn1] quit
[Sysname] bgp 100
[Sysname-bgp-default] address-family ipv4 unicast
```

New feature: Apply route update interval setting to withdrawn routes

Apply route update interval setting to withdrawn routes

About apply route update interval setting to withdrawn routes

Perform this task to apply the route update interval setting to withdrawn routes. BGP will send withdrawn and update messages at the specified intervals.

Procedure

1. Enter system view.
system-view
2. Enter BGP instance view.
bgp as-number [instance instance-name]
3. Apply route update interval setting to withdrawn routes.
route-update-interval withdrawn enable

By default, route update interval setting does not apply to withdrawn routes. BGP sends withdrawal messages for withdrawn routes immediately.

Command reference

route-update-interval withdrawn enable

Use **route-update-interval withdrawn enable** to apply route update interval setting to withdrawn routes.

Use **undo route-update-interval withdrawn enable** to restore the default.

Syntax

```
route-update-interval withdrawn enable
undo route-update-interval withdrawn enable
```

Default

Route update interval setting does not apply to withdrawn routes. BGP sends withdrawal messages for withdrawn routes immediately.

Views

BGP instance view

Predefined user roles

network-admin

Usage guidelines

This command does not take effect on routes that exist before the command is configured.

By default, for a route to be withdrawn, a BGP router immediately sends a withdrawal message to its peers. If the route changes frequently, the BGP router sends many updates for the route, resulting in routing flaps. To avoid this issue, use this command to apply the route update interval configured in

the **peer route-update-interval** command to withdraw routes. BGP will send withdrawn and update messages at the specified intervals.

Examples

In BGP instance view, apply route update interval setting to withdrawn routes.

```
<Sysname> system-view
```

```
[Sysname] bgp 100
```

```
[Sysname-bgp-default] route-update-interval withdrawn enable
```

New feature: Enabling prioritized withdrawal of specific routes

Enabling prioritized withdrawal of specific routes

About enabling prioritized withdrawal of specific routes

Perform this task to configure BGP to send the withdrawal messages of specific routes prior to other routes. This can achieve fast route switchover and reduce the traffic interruption time.

Procedure (IPv4 unicast/multicast address family)

1. Enter system view.
system-view
2. Enter BGP IPv4 unicast address family view, BGP-VPN IPv4 unicast address family view, or BGP IPv4 multicast address family view.
 - o Enter BGP IPv4 unicast address family view.
bgp as-number [instance instance-name]
address-family ipv4 [unicast]
 - o Enter BGP-VPN IPv4 unicast address family view.
bgp as-number [instance instance-name]
ip vpn-instance vpn-instance-name
address-family ipv4 [unicast]
 - o Enter BGP IPv4 multicast address family view.
bgp as-number [instance instance-name]
address-family ipv4 multicast
3. Enable BGP to send withdrawal messages of routes matching the specified routing policy prior to other routes.
update-first route-policy route-policy-name
By default, BGP does not send withdrawal messages of specific routes prior to other routes.

Procedure (IPv6 unicast/multicast address family)

1. Enter system view.
system-view
2. Enter BGP IPv6 unicast address family view, BGP-VPN IPv6 unicast address family view, or BGP IPv6 multicast address family view.
 - o Enter BGP IPv6 unicast address family view.
bgp as-number [instance instance-name]
address-family ipv6 [unicast]

- Enter BGP-VPN IPv6 unicast address family view.
`bgp as-number [instance instance-name]`
`ip vpn-instance vpn-instance-name`
`address-family ipv6 [unicast]`
 - Enter BGP IPv6 multicast address family view.
`bgp as-number [instance instance-name]`
`address-family ipv6 multicast`
3. Enable BGP to send withdrawal messages of routes matching the specified routing policy prior to other routes.
`update-first route-policy route-policy-name`
 By default, BGP does not send withdrawal messages of specific routes prior to other routes.

Command reference

update-first route-policy

Use `update-first route-policy` to configure BGP to send withdrawal messages of routes matching the specified routing policy prior to other routes.

Use `undo update-first route-policy` to restore the default.

Syntax

```
update-first route-policy route-policy-name
undo update-first route-policy
```

Default

BGP does not send withdrawal messages of any routes prior to other routes.

Views

BGP IPv4 unicast address family view
 BGP-VPN IPv4 unicast address family view
 BGP VPNv4 address family view
 BGP IPv6 unicast address family view
 BGP-VPN IPv6 unicast address family view
 BGP VPNv6 address family view
 BGP IPv4 multicast address family view
 BGP IPv6 multicast address family view

Predefined user roles

network-admin

Parameters

route-policy-name: Specifies a routing policy by its name, a case-sensitive string of 1 to 63 characters.

Usage guidelines

This command enables BGP to send the withdrawal messages of specific routes prior to other routes. This can achieve fast route switchover and reduce the traffic interruption time.

Examples

In BGP IPv4 unicast address family view, configure BGP to send withdrawal messages of routes matching routing policy **test-policy** prior to other routes

```
<Sysname> system-view
[Sysname] bgp 1
[Sysname-bgp-default] address-family ipv4
[Sysname-bgp-default-ipv4] update-first route-policy test-policy
```

New feature: Configuring BGP RPKI

Configuring BGP RPKI

About BGP RPKI

The AS_PATH attribute identifies the ASs through which a route has passed, and the AS that originated the route is the origin AS of the route. If the origin AS number of a route is incorrect, traffic transmission failure or even network collapse might occur.

To avoid this problem, you can configure the BGP Resource Public Key Infrastructure (RPKI) feature. It enables BGP to validate the origin AS of a route and determine whether to use and advertise the route based on the validation state.

Configuring RPKI connection parameters

About RPKI connections

A router establishes a TCP connection with an RPKI server to obtain the Route Origin Authorization (ROA) information used for RPKI validation. The router checks the connection to the RPKI server at the specified interval. If the router does not receive a response from the RPKI server within the specified time period, it tears down the connection to the RPKI server.

When the connection between a router and an RPKI server goes down (except when the **shutdown** command is executed), the router takes the following actions:

- Attempts to reconnect to the server.
- Places the ROA information obtained from the server in aging state, and starts the aging timer for the ROA information.

If the router reconnects to the server before the aging timer expires, it releases the ROA information from the aging state. If the router fails to reconnect to the server when the aging timer expires, it deletes the ROA information obtained from the server.

Restrictions and guidelines

Follow these restrictions and guidelines when you configure RPKI connection parameters:

- As a best practice, set an ROA information aging time longer than the time to wait for the response from the RPKI server.
- To tear down the connection to an RPKI server, execute the **undo port** command in RPKI server view.
- If you execute the **undo rpki** command, all configurations in RPKI view are removed.

Procedure

1. Enter system view.
system-view
2. Enter BGP instance view.

- bgp** *as-number* [**instance** *instance-name*]
- Enter BGP RPKI view.
rpki
 - Specify an RPKI server by its IP address and enter RPKI server view.
server [**vpn-instance** *vpn-instance-name*] **tcp** { *ipv4-address* | *ipv6-address* }
By default, no RPKI server is specified.
 - Specify the port number of the RPKI server.
port *port-number*
By default, the port number of the RPKI server is not specified.
To enable a router to establish a TCP connection with the RPKI server, you must configure the port number of the RPKI server on the router.
 - (Optional.) Specify the MD5 authentication password.
passwords { **cipher** | **simple** } *string*
By default, the RPKI server does not perform MD5 authentication.
The MD5 authentication password must be the same as the authentication password configured on the RPKI server.
MD5 authentication ensures the validity of the RPKI server and the security of BGP RPKI packets.
 - (Optional.) Set the RPKI connection check interval.
refresh-time *refresh-time*
By default, the RPKI connection check interval is 600 seconds.
 - (Optional.) Set the time to wait for the response from the RPKI server.
response-time *response-time*
By default, the time to wait for the response from the RPKI server is 30 seconds.
 - (Optional.) Set the aging time for the ROA information.
purge-time *purge-time*
By default, the aging time for the ROA information is 60 seconds.

Enabling BGP RPKI validation

About BGP RPKI validation

After you configure this feature, BGP validates the prefix and origin AS number of a received route and places the route to one of the following validation states:

- **Not-found**—No ROA matches the prefix.
- **Valid**—One or multiple ROAs match both the prefix and origin AS number.
- **Invalid**—One or multiple ROAs match the prefix, but none of the ROAs matches the origin AS number.

Restrictions and guidelines

If you configure this feature, BGP uses the local RPKI validation states. If you do not configure this feature, BGP uses the validation states in the received BGP routes.

Procedure

- Enter system view.
system-view
- Enter BGP instance view.

```
bgp as-number [ instance instance-name ]
```

3. Enter BGP RPKI view.

```
rpki
```

4. Enable BGP RPKI validation.

```
check-origin-validation
```

By default, BGP RPKI validation is disabled.

Applying the BGP RPKI validation state to optimal route selection

About applying the BGP RPKI validation state to optimal route selection

If multiple routes to the same destination are available, BGP first discards routes with unreachable next hops, and then selects the optimal route according to the following rules:

- Routes with a BGP RPKI validation state of **Valid** takes precedence over routes with a validation state of **Not-found** or **Invalid**.
- Routes with a BGP RPKI validation state of **Not-found** takes precedence over routes with a validation state of **Invalid**.
- Routes without a BGP RPKI validation state have the same priority as routes with a BGP RPKI validation state of **Not-found**.
- For routes that have the same BGP RPKI validation state, BGP selects the optimal route according to the BGP path selection rules.

You can configure a routing policy to filter routes based on the BGP RPKI validation state. For more information about routing policies, see "Configuring routing policies."

Procedure (IPv4 unicast)

1. Enter system view.

```
system-view
```

2. Enter BGP IPv4 unicast address family view or BGP-VPN IPv4 unicast address family view.

- o Enter BGP IPv4 unicast address family view.

```
bgp as-number [ instance instance-name ]
```

```
address-family ipv4 [ unicast ]
```

- o Enter BGP-VPN IPv4 unicast address family view.

```
bgp as-number [ instance instance-name ]
```

```
ip vpn-instance vpn-instance-name
```

```
address-family ipv4 [ unicast ]
```

3. Apply the BGP RPKI validation state to optimal route selection.

```
bestroute origin-as-validation [ allow-invalid ]
```

By default, BGP ignores the BGP RPKI validation state during optimal route selection.

To allow routes with a validation state of **Invalid** to participate in optimal route selection, you must specify the **allow-invalid** keyword.

Procedure (IPv6 unicast)

1. Enter system view.

```
system-view
```

2. Enter BGP IPv6 unicast address family view or BGP-VPN IPv6 unicast address family view.

- o Enter BGP IPv6 unicast address family view.

```
bgp as-number [ instance instance-name ]
```

- address-family ipv6 [unicast]**
 - Enter BGP-VPN IPv6 unicast address family view.
 - bgp as-number [instance instance-name]**
 - ip vpn-instance vpn-instance-name**
 - address-family ipv6 [unicast]**
- 3. Apply the BGP RPKI validation state to optimal route selection.
 - bestroute origin-as-validation [allow-invalid]**

By default, BGP ignores the BGP RPKI validation state during optimal route selection.

To allow routes with a validation state of **Invalid** to participate in optimal route selection, you must specify the **allow-invalid** keyword.

Advertising BGP RPKI validation state to a peer or peer group

Restrictions and guidelines

BGP advertises the BGP RPKI validation state to a peer or peer group through the extended community attribute. To enable this feature, you must first enable BGP to advertise the extended community attribute to the peer or peer group and make sure RPKI settings are correct.

In the current software version, BGP can advertise the BGP RPKI validation state only to IBGP peers and peer groups.

Procedure (IPv4 unicast)

1. Enter system view.
 - system-view**
2. Enter BGP IPv4 unicast address family view or BGP-VPN IPv4 unicast address family view.
 - Enter BGP IPv4 unicast address family view.
 - bgp as-number [instance instance-name]**
 - address-family ipv4 [unicast]**
 - Enter BGP-VPN IPv4 unicast address family view.
 - bgp as-number [instance instance-name]**
 - ip vpn-instance vpn-instance-name**
 - address-family ipv4 [unicast]**
3. Advertise the extended community attribute to a peer or peer group.
 - peer { group-name | ipv4-address [mask-length] }**
 - advertise-ext-community**

By default, BGP does not advertise the extended community attribute.
4. Advertise the BGP RPKI validation state to the specified peer or peer group.
 - peer { group-name | ipv4-address [mask-length] } advertise origin-as-validation**

By default, BGP does not advertise the BGP RPKI validation state.

Procedure (IPv6 unicast)

1. Enter system view.
 - system-view**
2. Enter BGP IPv6 unicast address family view or BGP-VPN IPv6 unicast address family view.
 - Enter BGP IPv6 unicast address family view.
 - bgp as-number [instance instance-name]**

- ```

 address-family ipv6 [unicast]

```
- Enter BGP-VPN IPv6 unicast address family view.

```

 bgp as-number [instance instance-name]
 ip vpn-instance vpn-instance-name
 address-family ipv6 [unicast]

```
  - 3. Advertise the extended community attribute to a peer or peer group.

```

peer { group-name | ipv6-address [prefix-length] }
advertise-ext-community

```

By default, BGP does not advertise the extended community attribute.
  - 4. Advertise the BGP RPKI validation state to the specified peer or peer group.

```

peer { group-name | ipv6-address [prefix-length] } advertise
origin-as-validation

```

By default, BGP does not advertise the BGP RPKI validation state.

## Resetting BGP RPKI sessions

Execute the following command in user view.

```

reset bgp [instance instance-name] rpki server [vpn-instance
vpn-instance-name] tcp { ipv4 address | ipv6 address }

```

## Display and maintenance commands for BGP security features

Execute **display** commands in any view.

### Displaying BGP (IPv4 unicast address family)

| Step                                                    | Command                                                                                                                                                              |
|---------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Display information about connections to RPKI servers.  | <b>display bgp</b> [ <b>instance</b> <i>instance-name</i> ] <b>rpki server</b> [ [ <b>vpn-instance</b> <i>vpn-instance-name</i> ] <i>ipv4-address</i> ]              |
| Display the ROA information obtained from RPKI servers. | <b>display bgp</b> [ <b>instance</b> <i>instance-name</i> ] <b>rpki table ipv4</b> [ <i>ipv4-address</i> <b>min</b> <i>min-length</i> <b>max</b> <i>max-length</i> ] |

### Displaying BGP (IPv6 unicast address family)

| Step                                                    | Command                                                                                                                                                              |
|---------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Display information about connections to RPKI servers.  | <b>display bgp</b> [ <b>instance</b> <i>instance-name</i> ] <b>rpki server</b> [ [ <b>vpn-instance</b> <i>vpn-instance-name</i> ] <i>ipv6-address</i> ]              |
| Display the ROA information obtained from RPKI servers. | <b>display bgp</b> [ <b>instance</b> <i>instance-name</i> ] <b>rpki table ipv6</b> [ <i>ipv6-address</i> <b>min</b> <i>min-length</i> <b>max</b> <i>max-length</i> ] |

# Command reference

## rpki

Use **rpki** to enter BGP RPKI view.

Use **undo rpki** to remove all configurations in BGP RPKI view.

### Syntax

```
rpki
undo rpki
```

### Views

BGP view

### Predefined user roles

network-admin

### Examples

```
Enter BGP RPKI view.
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] rpki
[Sysname-bgp-default-rpki]
```

## server tcp

Use **server tcp** to specify an RPKI server and enter BGP RPKI server view.

Use **undo server tcp** to cancel the configuration.

### Syntax

```
server [vpn-instance vpn-instance-name] tcp { ipv4-address |
ipv6-address }
undo server [vpn-instance vpn-instance-name] tcp { ipv4-address |
ipv6-address }
```

### Default

No RPKI server is specified.

### Views

BGP RPKI view

### Predefined user roles

network-admin

### Parameters

**vpn-instance** *vpn-instance-name*: Specifies a VPN instance by its name, a case-sensitive string of 1 to 31 characters. If you do not specify this option, the command specifies an RPKI server in the public network.

*ipv4-address*: Specifies an RPKI server by its IPv4 address.

*ipv6-address*: Specifies an RPKI server by its IPv6 address.

## Usage guidelines

You can execute this command multiple times to configure a router to establish TCP connections to multiple RPKI servers.

After you execute the **undo server** command, all configurations in BGP RPKI server view are deleted.

## Examples

# Specify an RPKI server with the IP address 1.1.1.1 and enter BGP RPKI server view.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] rpki
[Sysname-bgp-default-rpki] server tcp 1.1.1.1
[Sysname-bgp-default-rpki-server]
```

## port

Use **port** to specify the port number of the RPKI server.

Use **undo port** to restore the default.

## Syntax

```
port port-number
undo port
```

## Default

The port number of the RPKI server is not specified.

## Views

BGP RPKI server view

## Predefined user roles

network-admin

## Parameters

*port-number*: Specifies the port number of the RPKI server, in the range of 1 to 65535.

## Usage guidelines

After you specify the IP address and port number of the RPKI server, the device automatically establishes a TCP connection to the server.

Make sure the specified port number is the same as the port number used by the RPKI server.

Changing the port number will cause temporary connection interruption.

## Examples

# In BGP RPKI server view, specify 1234 as the port number of the RPKI server.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] rpki
[Sysname-bgp-default-rpki] server tcp 1.1.1.1
[Sysname-bgp-default-rpki-server] port 1234
```

## passwords

Use **passwords** to set the MD5 authentication password.

Use **undo passwords** to restore the default.

### Syntax

```
passwords { cipher | simple } string
undo passwords
```

### Default

The MD5 authentication password is not set.

### Views

BGP RPKI server view

### Predefined user roles

network-admin

### Parameters

**cipher**: Specifies a password in encrypted form.

**simple**: Specifies a password in plaintext form. For security purposes, the password specified in plaintext form will be stored in encrypted form.

*string*: Specifies the password. Its encrypted form is a case-sensitive string of 33 to 137 characters. Its plaintext form is a case-sensitive string of 1 to 80 characters.

### Usage guidelines

This command can enhance security in the following ways:

- Perform MD5 authentication when establishing TCP connections. Only the two parties that have the same password configured can establish TCP connections.
- Perform MD5 calculation on the packets transmitted over the TCP connections to avoid packet modification.

### Examples

```
In BGP RPKI server view, set the MD5 authentication password.
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] rpki
[Sysname-bgp-default-rpki] server tcp 1.1.1.1
[Sysname-bgp-default-rpki-server] passwords simple 123456
```

## refresh-time

Use **refresh-time** to set the RPKI connection check interval.

Use **undo refresh-time** to restore the default.

### Syntax

```
refresh-time refresh-time
undo refresh-time
```

### Default

The RPKI connection check interval is 600 seconds.

## Views

BGP RPKI server view

## Predefined user roles

network-admin

## Parameters

*refresh-time*: Specifies the RPKI connection check interval in the range of 15 to 3600 seconds.

## Usage guidelines

After a router establishes a TCP connection with an RPKI server, the router checks the connection at the specified interval. If the router does not receive a response from the RPKI server within the specified time period, it tears down the connection to the RPKI server.

## Examples

```
Set the RPKI connection check interval to 15 seconds.
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] rpki
[Sysname-bgp-default-rpki] server tcp 1.1.1.1
[Sysname-bgp-default-rpki-server] refresh-time 15
```

## response-time

Use **response-time** to set the time to wait for the response from the RPKI server.

Use **undo response-time** to restore the default.

## Syntax

```
response-time response-time
undo response-time
```

## Default

The time to wait for the response from the RPKI server is 30 seconds.

## Views

BGP RPKI server view

## Predefined user roles

network-admin

## Parameters

*response-time*: Specifies the time to wait for the response from the RPKI server, in the range of 15 to 3600 seconds.

## Usage guidelines

A router checks the connection to an RPKI server at the check interval. If the router does not receive a response from the RPKI server within the specified time period, it tears down the connection to the RPKI server.

## Examples

```
Set the time to wait for the response from RPKI server 1.1.1.1 to 15 seconds.
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] rpki
```

```
[Sysname-bgp-default-rpki] server tcp 1.1.1.1
[Sysname-bgp-default-rpki-server] response-time 15
```

## purge-time

Use **purge-time** to set the aging time for the ROA information.

Use **undo purge-time** to restore the default.

### Syntax

```
purge-time purge-time
undo purge-time
```

### Default

The aging time for the ROA information is 60 seconds.

### Views

BGP RPKI server view

### Predefined user roles

network-admin

### Parameters

*purge-time*: Specifies the aging time for the ROA information, in the range of 30 to 360 seconds.

### Usage guidelines

When the connection between a router and an RPKI server goes down (except when the **shutdown** command is executed), the router takes the following actions:

- Attempts to reconnect to the server.
- Places the ROA information obtained from the server in aging state, and starts the aging timer for the ROA information.

If the router reconnects to the server before the aging timer expires, it releases the ROA information from the aging state. If the router fails to reconnect to the server when the aging timer expires, it deletes the ROA information obtained from the server.

### Examples

```
Set the aging time for the ROA information to 150 seconds.
```

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] rpki
[Sysname-bgp-default-rpki] server tcp 1.1.1.1
[Sysname-bgp-default-rpki-server] purge-time 150
```

## check-origin-validation

Use **check-origin-validation** to enable BGP RPKI validation.

Use **undo check-origin-validation** to disable BGP RPKI validation.

### Syntax

```
check-origin-validation
undo check-origin-validation
```

## Default

BGP RPKI validation is disabled.

## Views

BGP RPKI view

## Predefined user roles

network-admin

## Usage guidelines

This command enables BGP to validate the prefix and origin AS number of a received route and place the route to one of the following validation states:

- **Not-found**—No ROA matches the prefix.
- **Valid**—One or multiple ROAs match both the prefix and origin AS number.
- **Invalid**—One or multiple ROAs match the prefix, but none of these ROAs matches the origin AS number.

You can configure a routing policy to filter routes based on the BGP RPKI validation state.

## Examples

```
Enable BGP RPKI validation.
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] rpki
[Sysname-bgp-default-rpki] check-origin-validation
```

## bestroute origin-as-validation

Use **bestroute origin-as-validation** to apply the BGP RPKI validation state to optimal route selection.

Use **undo bestroute origin-as-validation** to restore the default.

## Syntax

```
bestroute origin-as-validation [allow-invalid]
undo bestroute origin-as-validation
```

## Default

BGP ignores the BGP RPKI validation state during optimal route selection.

## Views

BGP IPv4 unicast address family view

BGP-VPN IPv4 unicast address family view

BGP IPv6 unicast address family view

BGP-VPN IPv6 unicast address family view

## Predefined user roles

network-admin

## Parameters

**allow-invalid**: Allows routes with a validation state of **Invalid** to participate in optimal route selection. If you do not specify this keyword, routes with a validation state of **Invalid** cannot participate in optimal route selection.

## Usage guidelines

If multiple routes to the same destination are available, BGP first discards routes with unreachable next hops, and then selects the optimal route according to the following rules:

- Routes with a BGP RPKI validation state of **Valid** takes precedence over routes with a validation state of **Not-found** or **Invalid**.
- Routes with a BGP RPKI validation state of **Not-found** takes precedence over routes with a validation state of **Invalid**.
- Routes without a BGP RPKI validation state have the same priority as routes with a BGP RPKI validation state of **Not-found**.

## Examples

# In BGP IPv4 unicast address family view, apply the BGP RPKI validation state to optimal route selection.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] address-family ipv4 unicast
[Sysname-bgp-default-ipv4] bestroute origin-as-validation
```

## peer advertise origin-as-validation

Use **peer advertise origin-as-validation** to configure BGP to advertise the BGP RPKI validation state to a peer or peer group.

Use **undo peer advertise origin-as-validation** to restore the default.

## Syntax

```
peer { group-name | ipv4-address [mask-length] | ipv6-address
[prefix-length] } advertise origin-as-validation

undo peer { group-name | ipv4-address [mask-length] | ipv6-address
[prefix-length] } advertise origin-as-validation
```

## Default

BGP does not advertise the BGP RPKI validation state to any peer or peer groups.

## Views

BGP IPv4 unicast address family view

BGP-VPN IPv4 unicast address family view

BGP IPv6 unicast address family view

BGP-VPN IPv6 unicast address family view

## Predefined user roles

network-admin

## Parameters

*group-name*: Specifies a peer group by its name, a case-sensitive string of 1 to 47 characters. The peer group must have been created.

*ipv4-address*: Specifies a peer by its IPv4 address. The peer must have been created.

*mask-length*: Specifies a mask length in the range of 0 to 32. You can use the *ipv4-address* and *mask-length* arguments together to specify a subnet. If you specify a subnet, this command specifies all dynamic peers in the subnet.

*ipv6-address*: Specifies a peer by its IPv6 address. The peer must have been created.

*prefix-length*: Specifies a prefix length in the range of 0 to 128. You can use the *ipv6-address* and *prefix-length* arguments together to specify a subnet. If you specify a subnet, this command specifies all dynamic peers in the subnet.

## Usage guidelines

BGP advertises the BGP RPKI validation state to a peer or peer group through the extended community attribute. To configure this command, you must first enable BGP to advertise the extended community attribute to the peer or peer group.

In the current software version, BGP can advertise the BGP RPKI validation state only to IBGP peers and peer groups.

## Examples

# In BGP IPv4 unicast address family view, configure BGP to advertise the BGP RPKI validation state to peer group **test**.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] address-family ipv4 unicast
[Sysname-bgp-default-ipv4] peer test advertise-ext-community
[Sysname-bgp-default-ipv4] peer test advertise origin-as-validation
```

## reset bgp rpki server

Use **reset bgp rpki server** to reset BGP RPKI sessions.

## Syntax

```
reset bgp [instance instance-name] rpki server [vpn-instance vpn-instance-name] tcp { ipv4 address | ipv6 address }
```

## Views

User view

## Predefined user roles

network-admin

## Parameters

**instance** *instance-name*: Specifies a BGP instance by its name, a case-sensitive string of 1 to 31 characters. If you do not specify this option, the command resets BGP RPKI sessions for the default BGP instance.

**vpn-instance** *vpn-instance-name*: Specifies an MPLS L3VPN instance by its name, a case-sensitive string of 1 to 31 characters. If you do not specify this option, the command resets BGP RPKI sessions for the public network.

*ipv4-address*: Specifies an RPKI server by its IPv4 address.

*ipv6-address*: Specifies an RPKI server by its IPv6 address.

## Usage guidelines

This command will cause temporary session interruption.

## Examples

# Reset the BGP RPKI session to RPKI server 2.2.2.1.

```
<Sysname> reset bgp rpki server tcp 2.2.2.1
```

## display bgp rpki server

Use **display bgp rpki server** to display information about connections to RPKI servers.

### Syntax

```
display bgp [instance instance-name] rpki server [[vpn-instance vpn-instance-name] ipv4-address]
```

```
display bgp [instance instance-name] rpki server [[vpn-instance vpn-instance-name] ipv6-address]
```

### Views

Any view

### Predefined user roles

network-admin

network-operator

### Parameters

**instance** *instance-name*: Specifies a BGP instance by its name, a case-sensitive string of 1 to 31 characters. If you do not specify this option, the command displays information for the default BGP instance.

**vpn-instance** *vpn-instance-name*: Specifies an MPLS L3VPN instance by its name, a case-sensitive string of 1 to 31 characters. If you do not specify this option, the command displays information for the public network.

*ipv4-address*: Specifies an RPKI server by its IPv4 address.

*ipv6-address*: Specifies an RPKI server by its IPv6 address.

### Examples

# Display brief information about connections to RPKI servers.

```
<Sysname> display bgp rpki server
```

| Server  | VPN-index | Port | State     | Time     | ROAs (IPv4/IPv6) |
|---------|-----------|------|-----------|----------|------------------|
| 1.1.1.2 | 0         | 1234 | Establish | 00:05:51 | 1/0              |
| 2.2.2.2 | 0         | 1234 | Establish | 00:06:07 | 3/1              |

**Table 33 Command output**

| Field           | Description                                                                                                                                                                                                                                                                                                |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Server          | IP address of the RPKI server.                                                                                                                                                                                                                                                                             |
| VPN-index       | VPN index.                                                                                                                                                                                                                                                                                                 |
| Port            | Port number of the RPKI server.                                                                                                                                                                                                                                                                            |
| State           | Connection state: <ul style="list-style-type: none"><li>• <b>Establish</b>—BGP has established a connection to the RPKI server.</li><li>• <b>Connect</b>—BGP is establishing a connection to the RPKI server.</li><li>• <b>Shutdown</b>—BGP has not established a connection to the RPKI server.</li></ul> |
| Time            | Duration of the current connection state.                                                                                                                                                                                                                                                                  |
| ROAs(IPv4/IPv6) | Number of IPv4/IPv6 ROAs.                                                                                                                                                                                                                                                                                  |

# Display detailed information about the connection to RPKI server 2.2.2.1.

```
<Sysname> display bgp rpki server 2.2.2.1
RPKI Cache-Server 2.2.2.1
 Port: TCP port 1234
 Local addr: 2.2.2.2, Local port: 14342
 Connect state: Establish
 Total byte Rx: 72
 Total byte Tx: 8
 Session ID: 1
 Serial number: 1
 Last PDU type 7, Time: 00:00:15
 Last disconnect reason: Response timer expired
```

**Table 34 Command output**

| Field                  | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RPKI Cache-Server      | IP address of the RPKI server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Port                   | Port number of the RPKI server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Local addr             | Local IP address of the connection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Local port             | Local port number of the connection.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Connect state          | Connection state: <ul style="list-style-type: none"> <li>• <b>Establish</b>—BGP has established a connection to the RPKI server.</li> <li>• <b>Connect</b>—BGP is establishing a connection to the RPKI server.</li> <li>• <b>Shutdown</b>—BGP has not established a connection to the RPKI server.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Total byte Rx          | Number of received bytes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Total byte Tx          | Number of sent bytes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Session ID             | Session ID assigned by the RPKI server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Serial number          | Serial number assigned by the RPKI server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Last PDU Type          | Type of the most recently received PDU.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Time                   | Duration of the current connection state.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Last disconnect reason | Reason for the most recent connection interruption: <ul style="list-style-type: none"> <li>• <b>Configure reset</b>—The port number used to establish the connection was changed or the <b>reset bgp rpki server</b> command was executed.</li> <li>• <b>Receive error report PDU</b>—Received error report packets from the server.</li> <li>• <b>Response timer expired</b>—No response was received within the response time from the RPKI server.</li> <li>• <b>Receive error PDU</b>—Received error packets.</li> <li>• <b>TCP connect failed</b>—The TCP connection was down.</li> <li>• <b>Shutdown port</b>—No port number was specified.</li> <li>• <b>Not enough memory</b>—The memory was insufficient.</li> <li>• <b>Receive cache reset PDU</b>—Received reset packets from the RPKI server.</li> </ul> |

## display bgp rpki table

Use **display bgp rpki table** to display the ROA information obtained from RPKI servers.

### Syntax

```
display bgp [instance instance-name] rpki table ipv4 [ipv4-address min
min-length max max-length]
```

```
display bgp [instance instance-name] rpki table ipv6 [ipv6-address min
min-length max max-length]
```

### Views

Any view

### Predefined user roles

network-admin

network-operator

### Parameters

**instance** *instance-name*: Specifies a BGP instance by its name, a case-sensitive string of 1 to 31 characters. If you do not specify this option, the command displays information for the default BGP instance.

**ipv4**: Displays information about ROAs with IPv4 prefixes.

*ipv4-address*: Specifies an IPv4 prefix.

**ipv6**: Displays information about ROAs with IPv6 prefixes.

*ipv6-address*: Specifies an IPv6 prefix.

*min-length*: Specifies the minimum prefix length. The value range for this argument is 0 to 32 for IPv4 prefixes and 0 to 128 for IPv6 prefixes.

*max-length*: Specifies the maximum prefix length. The value range for this argument is 0 to 32 for IPv4 prefixes and 0 to 128 for IPv6 prefixes.

### Examples

# Display brief information about ROAs with IPv4 prefixes.

```
<Sysname> display bgp rpki table ipv4
```

Total number of entries: 4

Status codes: S - stale, U - used

| Network | Mask-range | Origin-AS | Server  | Status |
|---------|------------|-----------|---------|--------|
| 1.2.3.4 | 8-32       | 100       | 1.1.1.2 | U      |
| 5.2.3.4 | 8-32       | 100       | 2.2.2.2 | U      |
| 6.6.6.6 | 8-32       | 100       | 2.2.2.2 | U      |
| 7.7.7.7 | 8-32       | 20        | 2.2.2.2 | U      |

**Table 35 Command output**

| Field                   | Description                  |
|-------------------------|------------------------------|
| Total number of entries | Total number of ROAs.        |
| Network                 | Network address.             |
| Mask-range              | Mask or prefix length range. |

| Field  | Description                                                                                                                              |
|--------|------------------------------------------------------------------------------------------------------------------------------------------|
| Server | IP address of the RPKI server.                                                                                                           |
| Status | ROA state: <ul style="list-style-type: none"> <li><b>U</b>—The ROA is available.</li> <li><b>S</b>—The ROA is in aging state.</li> </ul> |

# Display detailed information about ROAs with IPv4 prefixes.

```
<Sysname> display bgp rpki table ipv4 5.2.3.4 min 8 max 32
```

```
RPKI ROA entry for 5.2.3.4/8-32
```

```
Origin-AS: 100 from 2.2.2.1, used
```

**Table 36 Command output**

| Field     | Description                                                                                                                                                                                                                                                                     |
|-----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Origin-AS | ROA information: <ul style="list-style-type: none"> <li>AS number.</li> <li>IP address of the RPKI server.</li> <li>ROA state: <ul style="list-style-type: none"> <li><b>used</b>—The ROA is available.</li> <li><b>stale</b>—The ROA is in aging state.</li> </ul> </li> </ul> |

## New feature: Configuring BGP policy accounting

### Configuring BGP policy accounting

#### About BGP policy accounting

BGP policy accounting provides the following functionalities:

- Classifies IP traffic by BGP path attributes such as NEXT HOP, COMMUNITY, and AS\_PATH.
- Assigns an index to each class of traffic.
- Collects statistics based on indexes.

BGP policy accounting supports the following traffic classification policies:

- Source IP address-based policy**—Used to collect statistics for traffic from a specific source.  
BGP searches routes to the source address, obtains the traffic index for the routes, and collects statistics for the traffic identified by the index.
- Destination IP address-based policy**—Used to collect statistics for traffic destined to a specific address.  
BGP searches routes to the destination address, obtains the traffic index for the routes, and collects statistics for the traffic identified by the index.

#### Restrictions and guidelines

This feature is incompatible with fast forwarding as follows:

- When the source IP address-based policy is applied, fast forwarding does not work.
- When the destination IP address-based policy is applied, fast forwarding does not process traffic with a traffic index.

#### Procedure

- Enter system view.

**system-view**

2. Set a traffic index:
  - a. Configure a routing policy, and assign an index to the matching BGP routes by using the **apply traffic-index** command. For more information about the **apply traffic-index** command, see *Layer 3—IP Routing Command Reference*.
  - b. Configure BGP to use the routing policy. The **network**, **import-route**, **aggregate**, **peer route-policy**, and **peer default-route-advertise** commands can use a routing policy.
3. Enter interface view.  
**interface** *interface-type interface-number*
4. Enable BGP policy accounting on the interface.  
**bgp-policy accounting** { **input** | **output** } \* [ **source** ]  
By default, BGP policy accounting is disabled.  
If you do not specify the **source** keyword, the destination IP address-based policy applies.

## Command reference

### bgp-policy accounting

Use **bgp-policy accounting** to enable BGP policy accounting on an interface.

Use **undo bgp-policy accounting** to disable BGP policy accounting on an interface.

#### Syntax

```
bgp-policy accounting { input | output } * [source]
undo bgp-policy accounting { input | output } * [source]
```

#### Default

BGP policy accounting is disabled on an interface.

#### Views

Interface view

#### Predefined user roles

network-admin

#### Parameters

**input**: Enables BGP policy accounting for incoming traffic.

**output**: Enables BGP policy accounting for outgoing traffic.

**source**: Classifies IP traffic based on the source IP address. If you do not specify this keyword, the command classifies IP traffic based on the destination IP address.

#### Usage guidelines

BGP policy accounting provides the following functionalities:

- Classifies IP traffic by BGP path attributes such as NEXT HOP, COMMUNITY, and AS\_PATH.
- Assigns an index to each class of traffic.
- Collects statistics based on indexes.

BGP policy accounting supports the following traffic classification policies:

- **Source IP address-based policy**—BGP searches routes to the source address, obtains the traffic index for the routes, and collects statistics for the traffic identified by the index.

- **Destination IP address-based policy**—BGP searches routes to the destination address, obtains the traffic index for the routes, and collects statistics for the traffic identified by the index.

## Examples

# Enable BGP policy accounting based on source IP address for traffic on VLAN-interface 10.

```
<Sysname> system-view
[Sysname] interface vlan-interface 10
[Sysname-Vlan-interface10] bgp-policy accounting input output source
```

## display bgp-policy statistics

Use **display bgp-policy statistics** to display BGP policy accounting information.

## Syntax

```
display bgp-policy { ip | ipv6 } statistics { input | output } [interface
interface-type interface-number]
```

## Views

Any view

## Predefined user roles

network-admin  
network-operator

## Parameters

**ip**: Displays BGP policy accounting information for IPv4 traffic.

**ipv6**: Displays BGP policy accounting information for IPv6 traffic.

**input**: Displays BGP policy accounting information for incoming traffic.

**output**: Displays BGP policy accounting information for outgoing traffic.

**interface interface-type interface-number**: Displays BGP policy accounting information on an interface specified by the interface type and number. If you do not specify an interface, this command displays global BGP policy accounting information.

## Examples

# Display global BGP policy accounting information for incoming IPv4 traffic.

```
<Sysname> display bgp-policy ip statistics input
Statistics for the inbound direction:
```

| Traffic index | Packets | Bytes |
|---------------|---------|-------|
| 1             | 0       | 0     |
| 2             | 0       | 0     |
| 3             | 0       | 0     |
| 4             | 0       | 0     |
| 5             | 0       | 0     |
| 6             | 0       | 0     |
| 7             | 0       | 0     |
| 8             | 0       | 0     |
| 9             | 0       | 0     |
| 10            | 0       | 0     |

# Display global BGP policy accounting information for incoming IPv6 traffic.

```
<Sysname> display bgp-policy ipv6 statistics input
```

Statistics for the inbound direction:

| Traffic index | Packets | Bytes |
|---------------|---------|-------|
| 1             | 0       | 0     |
| 2             | 0       | 0     |
| 3             | 0       | 0     |
| 4             | 0       | 0     |
| 5             | 0       | 0     |
| 6             | 0       | 0     |
| 7             | 0       | 0     |
| 8             | 0       | 0     |
| 9             | 0       | 0     |
| 10            | 0       | 0     |

# Display BGP policy accounting information for outgoing IPv4 traffic on VLAN-interface 10.

```
<Sysname> display bgp-policy ip statistics output interface vlan-interface 10
```

Statistics for Vlan-interfacel0 in the outbound direction:

| Traffic index | Packets | Bytes |
|---------------|---------|-------|
| 1             | 0       | 0     |
| 2             | 0       | 0     |
| 3             | 0       | 0     |
| 4             | 0       | 0     |
| 5             | 0       | 0     |
| 6             | 0       | 0     |
| 7             | 0       | 0     |
| 8             | 0       | 0     |
| 9             | 0       | 0     |
| 10            | 0       | 0     |

# Display BGP policy accounting information for outgoing IPv6 traffic on VLAN-interface 10.

```
<Sysname> display bgp-policy ipv6 statistics output interface vlan-interface 10
```

Statistics for Vlan-interfacel0 in the outbound direction:

| Traffic index | Packets | Bytes |
|---------------|---------|-------|
| 1             | 0       | 0     |
| 2             | 0       | 0     |
| 3             | 0       | 0     |
| 4             | 0       | 0     |
| 5             | 0       | 0     |
| 6             | 0       | 0     |
| 7             | 0       | 0     |
| 8             | 0       | 0     |
| 9             | 0       | 0     |
| 10            | 0       | 0     |

**Table 37 Command output**

| Field   | Description                                             |
|---------|---------------------------------------------------------|
| Packets | Number of packets for traffic with the specified index. |
| Bytes   | Number of bytes for traffic with the specified index.   |

# New feature: Configuring BGP to immediately send updates for routes that match an IPv6 prefix list

## Configuring BGP to immediately send updates for routes that match an IPv6 prefix list

From this release, you can configure BGP to immediately send updates for routes that match an IPv6 prefix list.

### Command reference

#### bgp update-delay on-startup ipv6-prefix-list

Use **bgp update-delay on-startup ipv6-prefix-list** to configure BGP to immediately send route updates for routes that match an IPv6 prefix list.

Use **undo bgp update-delay on-startup ipv6-prefix-list** to restore the default.

#### Syntax

```
bgp update-delay on-startup ipv6-prefix-list ipv6-prefix-list-name
undo bgp update-delay on-startup ipv6-prefix-list
```

#### Default

No IPv6 prefix list is specified to filter routes.

#### Views

BGP instance view

#### Predefined user roles

network-admin

#### Parameters

*ipv6-prefix-list-name*: Specifies an IPv6 prefix list by its name, a case-sensitive string of 1 to 63 characters.

#### Usage guidelines

After the **bgp update-delay on-startup** command is configured, BGP delays sending updates for all routes on reboot. For BGP to immediately send updates for the specified IPv6 routes, execute the **bgp update-delay on-startup ipv6-prefix-list** command.

#### Examples

# In BGP instance view, configure BGP to send updates 100 seconds after reboot and immediately send updates for IPv6 routes that match IPv6 prefix list **bbb**.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] bgp update-delay on-startup 100
[Sysname-bgp-default] bgp update-delay on-startup ipv6-prefix-list bbb
```

#### Related commands

**bgp update-delay on-startup**

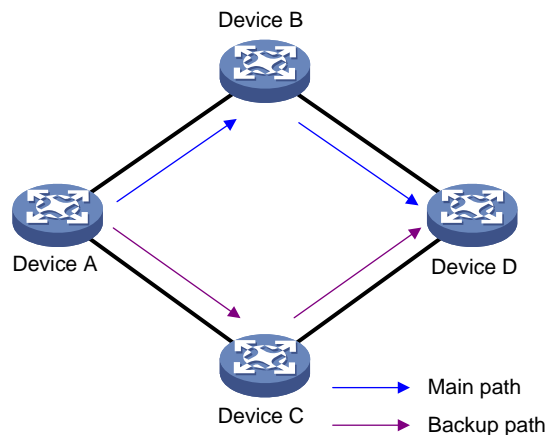
# New feature: BGP optimal route selection delay

## Configuring BGP optimal route selection delay

### About configuring BGP optimal route selection delay

As shown in Figure 6, the devices use BGP for communication. Device A has two paths to reach Device D. Typically Device A uses the main path to forward traffic to Device D. When the main path is unavailable, the backup path takes over. When the main path recovers, Device B sends a route update, which triggers an optimal route selection on Device A. Packet loss occurs if traffic is switched back to the main path when Device B has not updated its FIB table. To avoid this issue, you can perform this task on Device A to delay optimal route selection.

**Figure 6 Application scenario of optimal route selection delay**



### Restrictions and guidelines

Follow these restrictions and guidelines when you configure optimal route selection delay:

- The optimal route selection delay setting applies only when multiple effective routes with the same prefix exist after a route change occurs.
- For routes being delayed for optimal route selection, modifying the optimal route selection delay timer has the following effects:
  - If you modify the delay timer to a non-zero value, the routes are not affected, and they still use the original delay timer.
  - If you modify the delay timer to 0, the device performs optimal route selection immediately.
- If you configure the optimal route selection delay timer multiple times for an address family, the most recent configuration takes effect.
- The optimal route selection delay configuration does not apply to the following conditions:
  - A route change is caused by execution of a command or by route withdrawal.
  - After a route change occurs, only one route exists for a specific destination network.
  - An active/standby process switchover occurs.
  - A route change occurs among equal-cost routes.
  - Only the optimal and suboptimal routes exist when FRR is configured.
  - Optimal route selection is triggered by a redistributed route.

### Procedure (IPv4 unicast/multicast address family)

1. Enter system view.  
**system-view**

2. Enter BGP IPv4 unicast address family view, BGP-VPN IPv4 unicast address family view, or BGP IPv4 multicast address family view.
  - Enter BGP IPv4 unicast address family view.  
`bgp as-number [ instance instance-name ]`  
`address-family ipv4 [ unicast ]`
  - Enter BGP-VPN IPv4 unicast address family view.  
`bgp as-number [ instance instance-name ]`  
`ip vpn-instance vpn-instance-name`  
`address-family ipv4 [ unicast ]`
  - Enter BGP IPv4 multicast address family view.  
`bgp as-number [ instance instance-name ]`  
`address-family ipv4 multicast`
3. Set the optimal route selection delay timer.  
`route-select delay delay-value`  
 By default, the optimal route selection delay timer is 0 seconds, which means optimal route selection is not delayed.

#### Procedure (IPv6 unicast/multicast address family)

1. Enter system view.  
`system-view`
2. Enter BGP IPv6 unicast address family view, BGP-VPN IPv6 unicast address family view, or BGP IPv6 multicast address family view.
  - Enter BGP IPv6 unicast address family view.  
`bgp as-number [ instance instance-name ]`  
`address-family ipv6 [ unicast ]`
  - Enter BGP-VPN IPv6 unicast address family view.  
`bgp as-number [ instance instance-name ]`  
`ip vpn-instance vpn-instance-name`  
`address-family ipv6 [ unicast ]`
  - Enter BGP IPv6 multicast address family view.  
`bgp as-number [ instance instance-name ]`  
`address-family ipv6 multicast`
3. Set the optimal route selection delay timer.  
`route-select delay delay-value`  
 By default, the optimal route selection delay timer is 0 seconds, which means optimal route selection is not delayed.

## Command reference

### route-select delay

Use `route-select delay` to configure optimal route selection delay.

Use `undo route-select delay` to restore the default.

#### Syntax

`route-select delay delay-value`

`undo route-select delay`

## Default

The optimal route selection delay timer is 0 seconds, which means optimal route selection is not delayed.

## Views

BGP IPv4 unicast address family view

BGP-VPN IPv4 unicast address family view

BGP IPv6 unicast address family view

BGP-VPN IPv6 unicast address family view

BGP IPv4 multicast address family view

BGP IPv6 multicast address family view

BGP VPNv4 address family view

BGP VPNv6 address family view

BGP L2VPN address family view

BGP EVPN address family view

BGP IPv4 MDT address family view

BGP LS address family view

BGP IPv4 RT filter address family view

## Predefined user roles

network-admin

## Parameters

*delay-value*: Specifies the optimal route selection delay timer in the range of 0 to 600 seconds. A value of 0 means optimal route selection is not delayed.

## Usage guidelines

To avoid packet loss caused by path switchover, configure this command to delay optimal route selection upon a route change.

Follow these restrictions and guidelines when you configure the command:

- The optimal route selection delay setting applies only when multiple effective routes with the same prefix exist after a route change occurs.
- For routes being delayed for optimal route selection, modifying the optimal route selection delay timer has the following effects:
  - If you modify the delay timer to a non-zero value, the routes are not affected, and they still use the original delay timer.
  - If you modify the delay timer to 0, the device performs optimal route selection immediately.
- If you execute the command multiple times for an address family, the most recent configuration takes effect.
- The optimal route selection delay configuration does not apply to the following conditions:
  - A route change is caused by execution of a command or by route withdrawal.
  - After a route change occurs, only one route exists for a specific destination network.
  - An active/standby process switchover occurs.
  - A route change occurs among equal-cost routes.
  - Only the optimal and suboptimal routes exist when FRR is configured.

- Optimal route selection is triggered by a redistributed route.

## Examples

# In BGP IPv4 unicast address family view, set the optimal route selection delay timer to 100 seconds.

```
<Sysname> system-view
[Sysname] bgp 65009
[Sysname-bgp-default] address-family ipv4 unicast
[Sysname-bgp-default-ipv4] route-select delay 100
```

## Related commands

```
display bgp routing-table ipv4 multicast
display bgp routing-table ipv4 rtfilter
display bgp routing-table ipv4 unicast
display bgp routing-table ipv6 multicast
display bgp routing-table ipv6 unicast
```

# New feature: Specifying a VPN instance for the BMP server

## Specifying a VPN instance for the BMP server

From this release, you can specify a VPN instance for the BMP server.

## Command reference

### server vpn-instance

Use **server vpn-instance** to specify a VPN instance for the BMP server.

Use **undo server vpn-instance** to restore the default.

### Syntax

```
server vpn-instance vpn-instance-name
undo server vpn-instance
```

### Default

The BMP server has no VPN instance specified and resides in the public network.

### Views

BMP server view

### Predefined user roles

network-admin

### Parameters

*vpn-instance-name*: Specifies an MPLS L3VPN instance by its name, a case-sensitive string of 1 to 31 characters. The VPN instance must have been created.

## Usage guidelines

Deleting the specified VPN instance does not remove the configuration, but it disconnects the connection to the BMP server. After you create the VPN instance, the connection is automatically established again.

If you execute the command multiple times for a BMP server, the most recent configuration takes effect.

## Examples

# Specify VPN instance **vpna** for BMP server 5.

```
<Sysname> system-view
[Sysname] bmp server 5
[Sysname-bmpserver-5] server address 100.1.1.1 port 8888
[Sysname-bmpserver-5] server vpn-instance vpna
```

## Related commands

```
display bgp bmp server
ip vpn-instance (system view) (MPLS Command Reference)
server connect-interface
```

# New feature: Routing policy change delay

## Configuring routing policy change delay

### About configuring routing policy change delay

This feature makes a routing policy take effect after a delayed time interval, which prevents incomplete routing policy configuration from being issued to cause incorrect route advertisement.

Perform this task to set a routing policy change delay timer. Any of the following routing policy changes can trigger the delay timer, and does not take effect until the timer expires:

- Creating a routing policy.
- Adding, modifying, or deleting routing policy nodes, **if-match** clauses, and **apply** clauses.
- Adding, modifying, or deleting IPv4/IPv6 prefix lists, AS path lists, community lists, extended community lists, and MAC lists.
- Changing ACLs in the **if-match** clauses.

### Procedure

1. Enter system view.  
**system-view**
2. Set the routing policy change delay timer.  
**route-policy-change delay-time** { *time-value* | **unlimited** }  
By default, routing policy changes immediately take effect.

## Command reference

### route-policy-change delay-time

Use **route-policy-change delay-time** to set the routing policy change delay timer.

Use **undo route-policy-change delay-time** to restore the default.

## Syntax

```
route-policy-change delay-time { time-value | unlimited }
undo route-policy-change delay-time
```

## Default

Routing policy changes immediately take effect.

## Views

System view

## Predefined user roles

network-admin

## Parameters

*time-value*: Specifies the routing policy change delay timer in the range of 60 to 3600 seconds.

**unlimited**: Sets an infinite delay timer.

## Usage guidelines

This command makes a routing policy take effect after a delayed time interval, which prevents incomplete routing policy configuration from being issued to cause incorrect route advertisement.

Use this command to set a routing policy change delay timer. Any of the following routing policy changes can trigger the delay timer, and does not take effect until the timer expires:

- Creating a routing policy.
- Adding, modifying, or deleting routing policy nodes, **if-match** clauses, and **apply** clauses.
- Adding, modifying, or deleting IPv4/IPv6 prefix lists, AS path lists, community lists, extended community lists, and MAC lists.
- Changing ACLs in the **if-match** clauses.

To set a long delay time enough for you to complete routing policy configuration, you can specify the **unlimited** keyword for the command. Then, execute the **undo** form of the command after you complete the configuration.

Modifying the delay timer before it expires will cause the timer to restart.

## Examples

```
Set the routing policy change delay timer to 60 seconds.
<Sysname> system-view
[Sysname] route-policy-change delay-time 60
```

# New feature: Support of a routing policy for L3VNI configuration

## Specifying an L3VNI in if-match and apply clauses for a routing policy

From this release, you can specify an L3VNI in **if-match** and **apply** clauses for a routing policy.

## Command reference

### apply l3-vni

Use **apply l3-vni** to set an L3VNI for BGP EVPN routes.

Use **undo apply l3-vni** to restore the default.

## Syntax

**apply l3-vni** *vxlan-id*

**undo apply l3-vni**

## Default

No L3VNI is set.

## Views

Routing policy node view

## Predefined user roles

network-admin

## Parameters

*vxlan-id*: Specifies an L3VNI in the range of 0 to 16777215.

## Examples

# Configure node 10 in **permit** mode for routing policy **policy1** to set an L3VNI of 6 for BGP EVPN routes with an L3VNI of 8.

```
<Sysname> system-view
[Sysname] route-policy policy1 permit node 10
[Sysname-route-policy-policy1-10] if-match l3-vni 8
[Sysname-route-policy-policy1-10] apply l3-vni 6
```

## if-match l3-vni

Use **if-match l3-vni** to match BGP EVPN routes that have the specified L3VNI.

Use **undo if-match l3-vni** to restore the default.

## Syntax

**if-match l3-vni** *vxlan-id*

**undo if-match l3-vni**

## Default

No L3VNI match criterion is configured.

## Views

Routing policy node view

## Predefined user roles

network-admin

## Parameters

*vxlan-id*: Specifies an L3VNI in the range of 0 to 16777215.

## Examples

# Create node 10 in **permit** mode for routing policy **policy1** to match BGP EVPN routes that have an L3VNI of 8.

```
<Sysname> system-view
[Sysname] route-policy policy1 permit node 10
[Sysname-route-policy-policy1-10] if-match l3-vni 8
```

# New feature: Configuring the BGP RPKI validation state match criterion for a routing policy node

## Configuring the BGP RPKI validation state match criterion for a routing policy node

From this release, you can configure the BGP RPKI validation state match criterion for a routing policy node.

## Command reference

### if-match rpki

Use **if-match rpki** to set the BGP RPKI validation state match criterion.

Use **undo if-match rpki** to restore the default.

#### Syntax

```
if-match rpki { invalid | not-found | valid }
undo if-match rpki
```

#### Default

No BGP RPKI validation state match criterion is set.

#### Views

Routing policy node view

#### Predefined user roles

network-admin

#### Parameters

**invalid**: Matches BGP routes whose BGP RPKI validation state is **Invalid**.

**not-found**: Matches BGP routes whose BGP RPKI validation state is **Not found**.

**valid**: Matches BGP routes whose BGP RPKI validation state is **Valid**.

#### Examples

# Configure node 10 in **permit** mode for routing policy **policy1** to match BGP routes whose BGP RPKI validation state is **Valid**.

```
<Sysname> system-view
```

```
[Sysname] route-policy policy1 permit node 10
```

```
[Sysname-route-policy-policy1-10] if-match rpki valid
```

# New feature: Enabling logging for source MAC-based ARP attack detection

## Enabling logging for source MAC-based ARP attack detection

### About logging for source MAC-based ARP attack detection

When the device detects a source MAC-based ARP attack, it generates a log message and sends it to the information center. For information about the log destination and output rule configuration in the information center, see *Network Management and Monitoring Configuration Guide*.

The log messages help administrators to locate and solve problems.

### Procedure

1. Enter system view.  
**system-view**
2. Enable logging for source MAC-based ARP attack detection.  
**arp source-mac log enable**

By default, logging for source MAC-based ARP attack detection is disabled.

## Command reference

### arp source-mac log enable

Use **arp source-mac log enable** to enable logging for source MAC-based ARP attack detection.

Use **undo arp source-mac log enable** to disable logging for source MAC-based ARP attack detection.

### Syntax

```
arp source-mac log enable
undo arp source-mac log enable
```

### Default

Logging for source MAC-based ARP attack detection is disabled.

### Views

System view

### Predefined user roles

network-admin

### Usage guidelines

When the device detects a source MAC-based ARP attack, it generates a log message and sends it to the information center. For information about the log destination and output rule configuration in the information center, see *Network Management and Monitoring Configuration Guide*.

The log messages help administrators to locate and solve problems.

### Examples

```
Enable logging for source MAC-based ARP attack detection.
<Sysname> system-view
```

```
[Sysname] arp source-mac log enable
```

## New feature: Enabling ND attack detection logging

### Enabling ND attack detection logging

#### About ND attack detection logging

This feature allows a device to generate logs when it detects invalid ND packets. The log information helps administrators locate and solve problems. Each log records the following information:

- Attacked port numbers in a VLAN.
- IDs of the attacked Ethernet service instances in a VXLAN.
- Source IP address of the invalid ND packets.
- Source MAC address of the invalid ND packets.
- VLAN ID of the invalid ND packets.
- Total number of dropped ND packets.

#### Procedure

1. Enter system view.  
**system-view**
2. Enable ND attack detection logging.  
**ipv6 nd detection log enable**  
By default, ND attack detection logging is disabled.

### Command reference

#### ipv6 nd detection log enable

Use **ipv6 nd detection log enable** to enable ND attack detection logging.

Use **undo ipv6 nd detection log enable** to disable ND attack detection logging.

#### Syntax

```
ipv6 nd detection log enable
undo ipv6 nd detection log enable
```

#### Default

ND attack detection logging is disabled.

#### Views

System view

#### Predefined user roles

network-admin

#### Usage guidelines

This command allows a device to generate logs when it detects ND attacks. The log information helps administrators locate and solve problems. The ND attack detection logging feature sends the log message to the information center. The information center can then output log messages from different source modules to different destinations. For more information about information center, see *Network Management and Monitoring Configuration Guide*.

The device performance is degraded when the device outputs a large number of ND attack detection logs. You can disable ND attack detection logging to ensure the device performance.

## Examples

```
Enable ND attack detection logging.
<Sysname> system-view
[Sysname] ipv6 nd detection log enable
```

# New feature: Configuring BFD session flapping suppression

## Configuring BFD session flapping suppression

### About BFD session flapping suppression

When BFD detects a link failure, it tears down the BFD session and notifies the upper-layer protocol of the failure. When the upper-layer protocol re-establishes a neighbor relationship, the BFD session comes up again. BFD session flaps occur when a link fails and recovers repeatedly, which consumes significant system resources and causes network instability.

This feature allows you to suppress BFD session flapping by using the *initial-interval*, *secondary-interval*, and *maximum-interval* arguments.

- A BFD session is suppressed within the specified interval. The suppression time does not exceed the *maximum-interval*.
- After a BFD session goes down for the second time, it cannot be re-established within the *initial-interval*.
- After a BFD session goes down for the third time, it cannot be re-established within the *secondary-interval*.
- After a BFD session goes down for the fourth time and at any later time, the following rules apply:
  - If  $\text{secondary-interval} \times 2^{n-3}$  is smaller than or equal to the *maximum-interval*, the BFD session cannot be re-established within the  $\text{secondary-interval} \times 2^{n-3}$ .
  - If  $\text{secondary-interval} \times 2^{n-3}$  is greater than the *maximum-interval*, the BFD session cannot be re-established within the *maximum-interval*.

The letter n, starting from 4, is the number of times the BFD session flaps.

### Procedure

1. Enter system view.  
**system-view**
2. Configure BFD session flapping suppression.  
**bfd dampening [ maximum *maximum-interval* initial *initial-interval* secondary *secondary-interval* ]**

By default, BFD sessions are not suppressed.

The values for the *initial-interval* and *secondary-interval* arguments cannot be greater than the value for the *maximum-interval* argument.

# Command reference

## bfd dampening

Use **bfd dampening** to configure BFD session flapping suppression.

Use **undo bfd dampening** to disable BFD session flapping suppression.

### Syntax

```
bfd dampening [maximum maximum-interval initial initial-interval
secondary secondary-interval]
```

```
undo bfd dampening
```

### Default

BFD sessions are not suppressed.

### Views

System view

### Predefined user roles

network-admin

### Parameters

*maximum-interval*: Specifies the maximum amount of time a BFD session is suppressed, in the range of 1 to 3600 seconds. The default is 20 seconds.

*initial-interval*: Specifies the amount of time a BFD session is suppressed for the first time. The value range for the *initial-interval* argument is 1 to 3600 seconds. The default is 6 seconds.

*secondary-interval*: Specifies the amount of time a BFD session is suppressed for the second time. The value range for the *secondary-interval* argument is 1 to 3600 seconds. The default is 10 seconds.

### Usage guidelines

When BFD detects a link failure, it tears down the BFD session and notifies the upper-layer protocol of the failure. When the upper-layer protocol re-establishes a neighbor relationship, the BFD session comes up again. BFD session flaps occur when a link fails and recovers repeatedly, which consumes significant system resources and causes network instability.

This command allows you to suppress BFD session flapping by using the *initial-interval*, *secondary-interval*, and *maximum-interval* arguments.

- A BFD session is suppressed within the specified interval. The suppression time does not exceed the *maximum-interval*.
- After a BFD session goes down for the second time, it cannot be re-established within the *initial-interval*.
- After a BFD session goes down for the third time, it cannot be re-established within the *secondary-interval*.
- After a BFD session goes down for the fourth time and at any later time, the following rules apply:
  - If  $\text{secondary-interval} \times 2^{n-3}$  is smaller than or equal to the *maximum-interval*, the BFD session cannot be re-established within the  $\text{secondary-interval} \times 2^{n-3}$ .
  - If  $\text{secondary-interval} \times 2^{n-3}$  is greater than the *maximum-interval*, the BFD session cannot be re-established within the *maximum-interval*.

The letter n, starting from 4, is the number of times the BFD session flaps.

## Examples

```
Enable BFD session flapping suppression, and set the maximum-interval,
initial-interval, and secondary-interval to 12 seconds, 4 seconds, and 8 seconds,
respectively.
```

```
<Sysname> system-view
```

```
[Sysname] bfd dampening maximum 12 initial 4 secondary 8
```

## New feature: Configuring the timer that delays reporting the first BFD session establishment failure to the data link layer

### Configuring the timer that delays reporting the first BFD session establishment failure to the data link layer

#### About configuring the delay timer

If the local end is configured with the **bfd detect-interface** command, the BFD session for detecting the local interface state fails to be established when the following conditions exist:

- The remote end is not configured with the **bfd detect-interface** command.
- The local and remote ends have mismatching BFD authentication settings.

Perform this task to enable BFD to report the session establishment failure to the data link layer and set the data link layer state of the interface to DOWN(BFD). This helps you rapidly identify the interfaces for which BFD sessions fail to be established.

#### Procedure

1. Enter system view.  
**system-view**
2. Enter interface view.  
**interface** *interface-type* *interface-number*
3. Configure the timer that delays reporting the first BFD session establishment failure to the data link layer.  
**bfd detect-interface first-fail-timer** *seconds*

By default, the first BFD session establishment failure is not reported to the data link layer.

## Command reference

### bfd detect-interface first-fail-timer

Use **bfd detect-interface first-fail-timer** to configure the timer that delays reporting the first BFD session establishment failure to the data link layer.

Use **undo bfd detect-interface first-fail-timer** to restore the default.

#### Syntax

```
bfd detect-interface first-fail-timer seconds
```

```
undo bfd detect-interface first-fail-timer
```

## Default

The first BFD session establishment failure is not reported to the data link layer.

## Views

Interface view

## Predefined user roles

network-admin

## Parameters

*seconds*: Specifies the timeout time that reports the first BFD session establishment failure to the data link layer. The value range for this argument is 1 to 10000 seconds.

## Usage guidelines

If the BFD session fails to be established when the timer expires, BFD reports the failure to the data link layer and sets the data link layer state of the interface to DOWN(BFD). This behavior rapidly identifies the interfaces for which BFD sessions fail to be established. In this case, the BFD session state is displayed as Down in the **display bfd session** command output. The line protocol state of the interface is displayed as DOWN(BFD) in the **display interface** command output.

If the local end is configured with the **bfd detect-interface** command, the BFD session for detecting the local interface state fails to be established when the following conditions exist:

- The remote end is not configured with the **bfd detect-interface** command.
- The local and remote ends have mismatching BFD authentication settings.

## Examples

```
Configure the timer that delays reporting the first BFD session establishment failure as 10 seconds.
<Sysname> system-view
[Sysname] interface vlan-interface 11
[Sysname-Vlan-interface11] bfd detect-interface first-fail-timer 10
```

## Related commands

**bfd detect-interface**

**display interface** (*Interface Command Reference*)

# New feature: Enabling special processing for BFD sessions

## Enabling special processing for BFD sessions

### About enabling special processing for BFD sessions

By default, after a BFD session is established for detecting the local interface state, the local end performs the following actions when the specific events occur:

- Upon receipt of a BFD packet with the State field as AdminDown, the local end sets the session to down state, but does not notify the session down event to the data link layer.
- Upon a local authentication information change, the local end sets the session to down state if authentication information inconsistency still persists after a period of time.
- When the local session is up, the local end examines authentication information in incoming BFD packets. If the authentication information does not match on the two ends, the BFD session is declared down.

Perform this task to enable the local end to perform the following actions when the specific events occur:

- Upon receipt of a BFD packet with the State field as AdminDown, the local end notifies a session down event to the data link layer. This action helps rapidly discover the interfaces on which BFD sessions are manually shut down.
- Upon a local authentication information change, the local end immediately sets the session to down state. This action helps rapidly discover the interfaces with authentication information changes.
- When the local session is up, the local end ignores authentication information inconsistency. If a large number of BFD sessions exist, this action helps reduce the impact on device performance.

## Procedure

1. Enter system view.  
**system-view**
2. Enter interface view.  
**interface** *interface-type* *interface-number*
3. Enable special processing for BFD sessions.  
**bfd detect-interface special-processing** [ **admin-down** | **authentication-change** | **session-up** ] \*  
By default, all types of special processing for BFD sessions are disabled.

## Command reference

### bfd detect-interface special-processing

Use **bfd detect-interface special-processing** to enable special processing for BFD sessions.

Use **undo bfd detect-interface special-processing** to disable special processing for BFD sessions.

#### Syntax

```
bfd detect-interface special-processing [admin-down |
authentication-change | session-up] *

undo bfd detect-interface special-processing [admin-down |
authentication-change | session-up] *
```

#### Default

All types of special processing for BFD sessions are disabled.

#### Views

Interface view

#### Predefined user roles

network-admin

#### Parameters

**admin-down:** Notifies a session down event to the data link layer upon receipt of a BFD packet with the State field as AdminDown. This keyword helps rapidly discover interfaces that BFD sessions are manually shut down. If you do not specify this keyword, the device sets the BFD session state to Down, but does not notify the session down event to the data link layer.

**authentication-change:** Immediately sets the session to down state upon a local authentication information change. This keyword helps rapidly discover interfaces with authentication information changes. If you do not specify this keyword, the device sets the session to down state if authentication information inconsistency still persists after a period of time.

**session-up:** Ignores authentication information inconsistency when the local session is up. If there is a large number of BFD sessions, examining authentication information consistency affects device performance. If you do not specify this keyword, the device examines authentication information in incoming BFD packets when the local session state is up. If the authentication information does not match on the two ends, the BFD session is declared down.

## Usage guidelines

If you do not specify any parameters, this command enables or disables all types of special processing.

## Examples

# Enable all types of special processing for BFD sessions on VLAN-interface 11.

```
<Sysname> system-view
[Sysname] interface vlan-interface 11
[Sysname-Vlan-interface11] bfd detect-interface special-processing admin-down
authentication-change session-up
```

# New feature: SBFD

## Configuring SBFD

### About SBFD

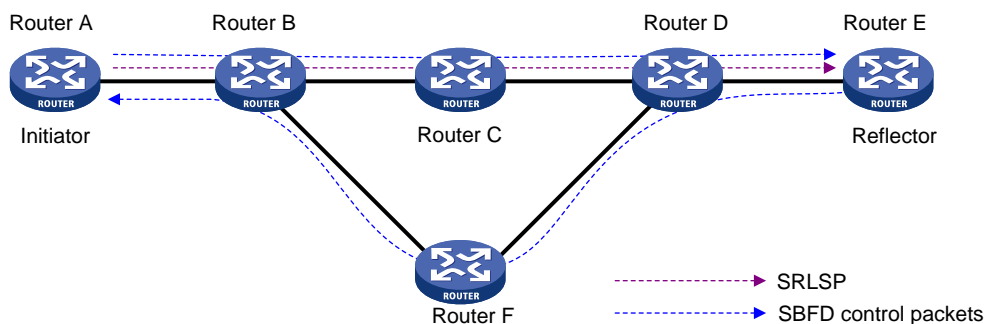
Seamless BFD (SBFD) is a unidirectional failure detection mechanism that provides shorter detection time than BFD. SBFD is used in scenarios where only one end of a link requires failure detection, such as MPLS TE tunneling and LSP establishment through LDP.

An SBFD session involves the following roles:

- **Initiator**—Periodically sends SBFD control packets to detect connectivity to remote entities. The initiator injects SBFD control packets into an MPLS TE tunnel or LDP LSP. The initiator initiates SBFD sessions and maintains SBFD session state.
- **Reflector**—Listens for incoming SBFD control packets on local entities and replies with response SBFD control packets.

As shown in [Figure 7](#), in an SR-based MPLS TE tunnel, the SRLSP from Router A to Router E is considered available if Router A (initiator) can receive response SBFD control packets from Router E (reflector).

**Figure 7 Initiator and reflector in SBFD**



## Restrictions and guidelines: SBFD configuration

A node can act as the initiator of one session and the reflector of another session at the same time.

An SBFD session can only be established statically. You must manually specify the remote discriminator for an SBFD session.

## Configuring the initiator

### Restrictions and guidelines

The remote discriminator in SBFD control packets sent by the initiator must be specified in the **sbfd local-discriminator** command on the reflector. If the discriminator is not specified on the reflector, the reflector does not reply with response SBFD control packets.

### Configuring the initiator for LDP LSP detection

1. Enter system view.

```
system-view
```

2. Enable BFD for MPLS.

```
mpls bfd enable
```

By default, BFD for MPLS is disabled.

For more information about this command, see MPLS OAM commands in *MPLS Command Reference*.

3. (Optional.) Set the detection time multiplier.

```
bfd multi-hop detect-multiplier value
```

The default setting is 5.

4. (Optional.) Set the minimum interval for sending SBFD control packets.

```
bfd multi-hop min-transmit-interval interval
```

The default value varies by device model.

5. Configure SBFD to verify LSP connectivity for an FEC.

```
mpls sbfd dest-addr mask-length [nexthop nexthop-address] remote
remote-id [template template-name]
```

By default, SBFD is not configured to verify LSP connectivity for an FEC.

For more information about this command, see MPLS OAM commands in *MPLS Command Reference*.

### Configuring the initiator for MPLS TE tunnel detection

1. Enter system view.

```
system-view
```

2. Enable BFD for MPLS.

```
mpls bfd enable
```

By default, BFD for MPLS is disabled.

For more information about this command, see MPLS OAM commands in *MPLS Command Reference*.

3. (Optional.) Set the detection time multiplier.

```
bfd multi-hop detect-multiplier value
```

The default setting is 5.

4. (Optional.) Set the minimum interval for sending SBFD control packets.

```
bfd multi-hop min-transmit-interval interval
```

The default value varies by device model.

5. Enter the view of an MPLS TE tunnel interface.

```
interface tunnel number [mode mpls-te]
```

6. Enable SBFD to verify MPLS TE tunnel connectivity.

```
mpls sbfd remote remote-id [template template-name]
```

By default, SBFD is not enabled for an MPLS TE tunnel.

For more information about this command, see MPLS OAM commands in *MPLS Command Reference*.

## Configuring the reflector

### Restrictions and guidelines

If you configure an IPv4 address as the local discriminator, the device automatically converts it to an integer. Configure an IPv4 address local discriminator only when it is required for interoperability with a third-party device.

The reflector replies with a response SBFD control packet only when the remote discriminator in the SBFD control packet sent from the initiator is specified in the **sbfd local-discriminator** command.

### Procedure

1. Enter system view.

```
system-view
```

2. Configure a local discriminator.

```
sbfd local-discriminator { ipv4-address | integer-value }
```

By default, no local discriminator is configured.

You can execute this command multiple times to configure multiple local discriminators.

## Command reference

### display sbfd session

Use **display sbfd session** to display SBFD session information.

#### Syntax

```
display sbfd session { initiator | reflector } [discriminator value | verbose]
```

#### Views

Any view

#### Predefined user roles

network-admin

network-operator

#### Parameters

**initiator**: Displays SBFD session information for the initiator.

**reflector**: Displays SBFD session information for the reflector.

**discriminator** *value*: Specifies a local discriminator in the range of 1 to 4294967295.

**verbose**: Displays detailed information about all SBFD sessions.

## Usage guidelines

If you do not specify the **discriminator** or **verbose** keyword, this command displays brief information about all SBFD sessions.

## Examples

# Display brief information about all SBFD sessions for the initiator.

```
<Sysname> display sbfd session initiator
```

```
Total Session Num: 1 Up Session Num: 1
```

```
SBFD Session (MPLS LSP):
```

```
Local Discr: 1 Remote Discr: 1000001
Source IP: 1.1.1.1
Destination IP: 127.0.0.1
Session State: Up
Hold Time: 2184ms
```

**Table 38 Command output**

| Field                   | Description                                                                                                                                               |
|-------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Total Session Num       | Total number of SBFD sessions.                                                                                                                            |
| Up Session Num          | Number of active SBFD sessions.                                                                                                                           |
| SBFD Session (MPLS LSP) | SBFD sessions created by MPLS.                                                                                                                            |
| Local Discr             | Local discriminator of the session.                                                                                                                       |
| Remote Discr            | Remote discriminator of the session.                                                                                                                      |
| Source IP               | Source IP address of the session.                                                                                                                         |
| Destination IP          | Destination IP address of the session.                                                                                                                    |
| Session State           | Session state: <b>Down</b> or <b>Up</b> .                                                                                                                 |
| Hold Time               | Length of time before the session detection timer expires, in milliseconds.<br>For an SBFD session in <b>Down</b> state, this field displays <b>0ms</b> . |

# Display detailed information about the SBFD session with local discriminator 2 for the initiator.

```
<Sysname> display sbfd session initiator discriminator 2
```

```
Local Discr: 2 Remote Discr: 1000002
Source IP: 1.1.1.2 Destination IP: 127.0.0.2
Session State: Up Hold Time: 4130ms
Min Tx Inter: 500ms Act Tx Inter: 1000ms
Detect Inter: 5000ms Running Up for: 00:01:30
Rx Count: 100 Tx Count: 100
Slot: 0
Protocol: MPLS_LSPV
Diag Info: No Diagnostic
Template Name: abc
```

**Table 39 Command output**

| Field          | Description                                                                                                                                                                                                                                                                                                                                     |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Local Discr    | Local discriminator of the session.                                                                                                                                                                                                                                                                                                             |
| Remote Discr   | Remote discriminator of the session.                                                                                                                                                                                                                                                                                                            |
| Source IP      | Source IP address of the session.                                                                                                                                                                                                                                                                                                               |
| Destination IP | Destination IP address of the session.                                                                                                                                                                                                                                                                                                          |
| Session State  | Session state: <b>Down</b> or <b>Up</b> .                                                                                                                                                                                                                                                                                                       |
| Hold Time      | Length of time before the session detection timer expires, in milliseconds.<br>For an SBFD session in <b>Down</b> state, this field displays <b>0ms</b> .                                                                                                                                                                                       |
| Min Tx Inter   | Minimum SBFD packet transmission interval in milliseconds.                                                                                                                                                                                                                                                                                      |
| Act Tx Inter   | Actual SBFD packet transmission interval in milliseconds.                                                                                                                                                                                                                                                                                       |
| Detect Inter   | Actual session detection timer in milliseconds.                                                                                                                                                                                                                                                                                                 |
| Running up for | Time period for which the session has been up.                                                                                                                                                                                                                                                                                                  |
| Rx Count       | Number of packets received.                                                                                                                                                                                                                                                                                                                     |
| Tx Count       | Number of packets sent.                                                                                                                                                                                                                                                                                                                         |
| Diag Info      | Diagnostic information about the session: <ul style="list-style-type: none"> <li>• <b>No Diagnostic.</b></li> <li>• <b>Control Detection Time Expired</b>—The SBFD session goes down because local detection times out.</li> <li>• <b>Neighbor Signaled Session Down</b>—The remote end notifies the local end of SBFD session down.</li> </ul> |
| Template Name  | Template associated with the session. If no template is associated, this field is not displayed.                                                                                                                                                                                                                                                |

# Display detailed information about all SBFD sessions for the initiator.

```
<Sysname> display sbfd session initiator verbose
```

```
Total Session Num: 1 Up Session Num: 1
```

```
SBFD Session (MPLS LSP):
```

```

 Local Discr: 1 Remote Discr: 1000001
 Source IP: 1.1.1.1 Destination IP: 127.0.0.1
 Session State: Up Hold Time: 4130ms
 Min Tx Inter: 500ms Act Tx Inter: 1000ms
 Detect Inter: 5000ms Running Up for: 00:01:30
 Rx Count: 100 Tx Count: 100
 Slot: 0
 Protocol: MPLS_LSPV
 Diag Info: No Diagnostic
 Template Name: abc
```

**Table 40 Command output**

| Field             | Description                     |
|-------------------|---------------------------------|
| Total Session Num | Total number of SBFD sessions.  |
| Up Session Num    | Number of active SBFD sessions. |

| Field                   | Description                                                                                                                                                                                                                                                                                                                                     |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SBFD Session (MPLS LSP) | SBFD sessions created by MPLS.                                                                                                                                                                                                                                                                                                                  |
| Local Discr             | Local discriminator of the session.                                                                                                                                                                                                                                                                                                             |
| Remote Discr            | Remote discriminator of the session.                                                                                                                                                                                                                                                                                                            |
| Source IP               | Source IP address of the session.                                                                                                                                                                                                                                                                                                               |
| Destination IP          | Destination IP address of the session.                                                                                                                                                                                                                                                                                                          |
| Session State           | Session state: <b>Down</b> or <b>Up</b> .                                                                                                                                                                                                                                                                                                       |
| Hold Time               | Length of time before the session detection timer expires, in milliseconds.<br>For an SBFD session in <b>Down</b> state, this field displays <b>0ms</b> .                                                                                                                                                                                       |
| Min Tx Inter            | Minimum SBFD packet transmission interval.                                                                                                                                                                                                                                                                                                      |
| Act Tx Inter            | Actual SBFD packet transmission interval.                                                                                                                                                                                                                                                                                                       |
| Detect Inter            | Actual session detection timer.                                                                                                                                                                                                                                                                                                                 |
| Running up for          | Time period for which the session has been up.                                                                                                                                                                                                                                                                                                  |
| Rx Count                | Number of packets received.                                                                                                                                                                                                                                                                                                                     |
| Tx Count                | Number of packets sent.                                                                                                                                                                                                                                                                                                                         |
| Diag Info               | Diagnostic information about the session: <ul style="list-style-type: none"> <li>• <b>No Diagnostic.</b></li> <li>• <b>Control Detection Time Expired</b>—The SBFD session goes down because local detection times out.</li> <li>• <b>Neighbor Signaled Session Down</b>—The remote end notifies the local end of SBFD session down.</li> </ul> |
| Template Name           | Template associated with the session. If no template is associated, this field is not displayed.                                                                                                                                                                                                                                                |

# Display brief information about all SBFD sessions for the reflector.

```
<Sysname> display sbfd session reflector
```

```
Local discriminator 1000001 on slot 0
Configured mode: Explicit value
Total Session Num: 1
```

```

Local Discr: 1000001 Remote Discr: 1
Source IP: 2.2.2.1
Destination IP: 1.1.1.1
Hold Time: 2200ms
VRF Name: vpn1
```

**Table 41 Command output**

| Field             | Description                                                                                                                                       |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| Configured mode   | Discriminator type: <ul style="list-style-type: none"> <li>• <b>IP Address</b>—IPv4 address.</li> <li>• <b>Explicit Value</b>—Integer.</li> </ul> |
| Total Session Num | Total number of SBFD sessions.                                                                                                                    |
| Local Discr       | Local discriminator of the session.                                                                                                               |

| Field          | Description                                                                                                                                                                |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Remote Discr   | Remote discriminator of the session.                                                                                                                                       |
| Source IP      | Source IP address of the session.                                                                                                                                          |
| Destination IP | Destination IP address of the session.                                                                                                                                     |
| Hold Time      | Length of time before the session detection timer expires, in milliseconds.<br>For an SBFD session in <b>Down</b> state, this field displays <b>0ms</b> .                  |
| VRF Name       | VPN instance bound to the interface where the packets of the SBFD session are received.<br>If no VPN instance is bound to the interface, this field displays a hyphen (-). |

# Display detailed information about the SBFD session with local discriminator 1000001 for the reflector.

```
<Sysname> display sbfd session reflector discriminator 1000001
```

```
Local discriminator 1000001 on slot 0
```

```
Configured mode: IP Address
```

```
Total Session Num: 2
```

|                      |                          |
|----------------------|--------------------------|
| Local Discr: 1000001 | Remote Discr: 1          |
| Source IP: 2.2.2.1   | Destination IP: 1.1.1.1  |
| Min Tx Inter: 500ms  | Min Rx Inter: 10ms       |
| Detect Multi: 5      | Running Up for: 00:02:30 |
| Rx Count: 100        | Tx Count: 100            |
| Max Rx Period: 498ms | Max Tx Period: 1280us    |
| Hold Time: 2200ms    |                          |
| VRF Name: vpn1       |                          |

|                      |                          |
|----------------------|--------------------------|
| Local Discr: 1000001 | Remote Discr: 2          |
| Source IP: 2.2.2.2   | Destination IP: 1.1.1.2  |
| Min Tx Inter: 1000ms | Min Rx Inter: 3790ms     |
| Detect Multi: 5      | Running Up for: 00:10:59 |
| Rx Count: 123        | Tx Count: 122            |
| Max Rx Period: 980ms | Max Tx Period: 580us     |
| Hold Time: 2500ms    |                          |
| VRF Name: vpn2       |                          |

**Table 42 Command output**

| Field             | Description                                                                                                                                   |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| Configured mode   | Discriminator type: <ul style="list-style-type: none"> <li><b>IP Address</b>—IPv4 address.</li> <li><b>Explicit Value</b>—Integer.</li> </ul> |
| Total Session Num | Total number of SBFD sessions.                                                                                                                |
| Local Discr       | Local discriminator of the session.                                                                                                           |
| Remote Discr      | Remote discriminator of the session.                                                                                                          |

| Field          | Description                                                                                                                                                                |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Source IP      | Source IP address of the session.                                                                                                                                          |
| Destination IP | Destination IP address of the session.                                                                                                                                     |
| Hold Time      | Length of time before the session detection timer expires, in milliseconds.<br>For an SBFD session in <b>Down</b> state, this field displays <b>0ms</b> .                  |
| Min Tx Inter   | Minimum transmission interval of incoming SBFD packets.                                                                                                                    |
| Min Rx Inter   | Minimum SBFD packet receiving interval.                                                                                                                                    |
| Detect Multi   | Session detection time multiplier carried in incoming SBFD packets.                                                                                                        |
| Running up for | Time period for which the session has been up.                                                                                                                             |
| Rx Count       | Number of packets received.                                                                                                                                                |
| Tx Count       | Number of packets sent.                                                                                                                                                    |
| Max Rx Period  | Maximum SBFD packet receiving interval in milliseconds.                                                                                                                    |
| Max Tx Period  | Maximum response packet transmission interval, in microseconds.                                                                                                            |
| Hold Time      | Length of time before the session detection timer expires, in milliseconds.<br>For an SBFD session in <b>Down</b> state, this field displays <b>0ms</b> .                  |
| VRF Name       | VPN instance bound to the interface where the packets of the SBFD session are received.<br>If no VPN instance is bound to the interface, this field displays a hyphen (-). |

## Related commands

**sbfd local-discriminator**

## sbfd local-discriminator

Use **sbfd local-discriminator** to set the local discriminator for the reflector.

Use **undo sbfd local-discriminator** to remove the local discriminator setting for the reflector.

## Syntax

```
sbfd local-discriminator { ipv4-address | integer-value }
undo sbfd local-discriminator { ipv4-address | integer-value }
```

## Default

No local discriminator is set for the reflector.

## Views

System view

## Predefined user roles

network-admin

## Parameters

*ipv4-address*: Specifies a valid unicast IPv4 address as the local discriminator.

*integer-value*: Specifies an integer as the local discriminator.

## Usage guidelines

In a BFD session, both ends periodically send BFD packets to detect bidirectional forwarding path failures. SBFD is a unidirectional failure detection mechanism that provides shorter detection time than BFD. SBFD is used in scenarios where only one end of a link requires failure detection, such as MPLS TE tunnel establishment through static SRLSP.

An SBFD session involves the following roles:

- **Initiator**—Periodically sends SBFD control packets to detect connectivity to remote entities. An initiator initiates SBFD sessions and maintains SBFD session state.
- **Reflector**—Listens for incoming SBFD control packets on local entities and replies with response SBFD control packets.

The remote discriminator in SBFD control packets sent by the initiator must be specified in the **sbfd local-discriminator** command. Otherwise, the reflector does not reply with response SBFD control packets.

You can execute this command multiple times to configure multiple local discriminators.

A node can act as the initiator of one session and the reflector of another session at the same time.

## Examples

# Specify 1.1.1.1 as the local discriminator for the reflector.

```
<Sysname> system-view
```

```
[Sysname] sbfd local-discriminator 1.1.1.1
```

# Specify 1000001 as the local discriminator for the reflector.

```
<Sysname> system-view
```

```
[Sysname] sbfd local-discriminator 1000001
```

## Related commands

```
display sbfd session
```

# New feature: Associating Track with a control-mode BFD session

## Associating Track with a control-mode BFD session

### About Track association with a control-mode BFD session

From this version, you can associate a track entry with a control-mode BFD session. For more information about BFD, see BFD configuration in *High Availability Configuration Guide*.

The associated Track and BFD operate as follows:

- If the BFD detects that the link fails, it informs the Track module of the link failure. The Track module sets the track entry to Negative state.
- If the BFD detects that the link is operating correctly, the Track module sets the track entry to Positive state.

## Procedure

1. Enter system view.

```
system-view
```

2. Create a track entry, and associate it with a control-mode BFD session.

```
track track-entry-number bfd ctrl [interface interface-type
interface-number | vpn-instance vpn-instance-name] remote ip
```

```
remote-ip-address local ip local-ip-address [delay { negative
negative-time | positive positive-time } *]
```

## Command reference

### track bfd ctrl

Use **track bfd ctrl** to create a track entry and associate it with a control-mode BFD session.

Use **undo track ctrl** to remove the track entry.

#### Syntax

```
track track-entry-number bfd ctrl [interface interface-type
interface-number | vpn-instance vpn-instance-name] remote ip
remote-ip-address local ip local-ip-address [delay { negative
negative-time | positive positive-time } *]

undo track track-entry-number
```

#### Default

No track entries exist.

#### Views

System view

#### Predefined user roles

network-admin

#### Parameters

***track-entry-number***: Specifies the track entry ID in the range of 1 to 1024.

***interface interface-type interface-number***: Specifies the BFD control packets' outgoing interface by its type and number. If you do not specify an outgoing interface, the outgoing interface found through FIB table lookup is used.

***vpn-instance vpn-instance-name***: Specifies an MPLS L3VPN instance by its name, a case-sensitive string of 1 to 31 characters. If you do not specify this option, the BFD session is on the public network.

***remote ip remote-ip-address***: Specifies the destination IP address of the BFD control packets. The specified IP address must be the IP address of a directly connected interface.

***local ip local-ip-address***: Specifies the source IP address of the BFD control packets. The specified IP address must be the IP address of a directly connected interface.

***delay***: Specifies the delay before the Track module notifies the application modules of the track entry state change. If you do not specify this keyword, the Track module notifies the application modules immediately when the track entry state changes.

***negative negative-time***: Specifies the delay before the Track module notifies the application modules that the track entry state has changed to Negative. The *negative-time* argument represents the notification delay in the range of 1 to 300 seconds.

***positive positive-time***: Specifies the delay before the Track module notifies the application modules that the track entry state has changed to Positive. The *positive-time* argument represents the notification delay in the range of 1 to 300 seconds.

## Usage guidelines

After a track entry is created, you can only use the **track bfd ctrl delay** command to modify its notification delay settings. If you execute the **track bfd ctrl delay** command multiple times, the most recent configuration takes effect.

To modify other settings for the track entry, execute the **undo track** command to remove the track entry, and then execute the **track bfd ctrl** command again.

When you associate Track with a BFD session, do not use the virtual IP address of a VRRP group as the local or remote address of the BFD session.

## Examples

# Associate track entry 1 with a control-mode BFD session. The BFD control packets use destination IP address 192.168.1.1, source IP address 192.168.1.2, and outgoing interface VLAN interface 2.

```
<Sysname> system-view
```

```
[Sysname] track 1 bfd ctrl interface vlan-interface 2 remote ip 192.168.1.1 local ip 192.168.1.2
```

## Related commands

**display track**

# New feature: Enabling the device to send IGMP general queries upon a path change

## Enabling the device to send IGMP general queries upon a path change

### About sending IGMP general queries upon a path change

Perform this task to enable the device to immediately send an IGMP general query to all ports in the affected VLAN or VSI when a path changes. This feature ensures that the device can update the member port information in multicast forwarding entries and switch multicast traffic to the new path promptly.

### Restrictions and guidelines

Perform this task on the IGMP snooping querier in a network of which the topology frequently changes.

### Procedure

1. Enter system view.  
**system-view**
2. Enter IGMP-snooping view.  
**igmp-snooping**
3. Enable the device to send IGMP general queries upon a path change.  
**send-query enable**

By default, the device does not send IGMP general queries upon a path change.

## Command reference

### send-query enable

Use **send-query enable** to enable the device to send IGMP general queries upon a path change.

Use **undo send-query enable** to restore the default.

#### Syntax

**send-query enable**

**undo send-query enable**

#### Default

The device does not send IGMP general queries upon a path change.

#### Views

IGMP-snooping view

#### Predefined user roles

network-admin

#### Usage guidelines

Configure this command to enable the device to immediately send an IGMP general query to all ports in the affected VLAN or VSI when the path changes. This feature ensures that the device can update the member port information in multicast forwarding entries and switch multicast traffic to the new path promptly.

#### Examples

```
Enable the device to send IGMP general queries upon a path change.
<Sysname> system-view
[Sysname] igmp-snooping
[Sysname-igmp-snooping] send-query enable
```

#### Related commands

**igmp-snooping**

## New feature: Enabling the device to send MLD general queries upon a path change

### Enabling the device to send MLD general queries upon a path change

#### About sending MLD general queries upon a path change

Perform this task to enable the device to immediately send an MLD general query to all ports in the affected VLAN or VSI when the path changes. This feature ensures that the device can update the member port information in IPv6 multicast forwarding entries and switch IPv6 multicast traffic to the new path promptly.

#### Restrictions and guidelines

Perform this task on the MLD snooping querier in a network of which the topology frequently changes.

## Procedure

1. Enter system view.  
**system-view**
  2. Enter MLD-snooping view.  
**mld-snooping**
  3. Enable the device to send MLD general queries upon a path change.  
**send-query enable**
- By default, the device does not send MLD general queries upon a path change.

## Command reference

### send-query enable

Use **send-query enable** to enable the device to send MLD general queries upon a path change.

Use **undo send-query enable** to restore the default.

#### Syntax

```
send-query enable
undo send-query enable
```

#### Default

The device does not send MLD general queries upon a path change.

#### Views

MLD-snooping view

#### Predefined user roles

network-admin

#### Usage guidelines

Perform this task to enable the device to immediately send an MLD general query to all ports in the affected VLAN or VSI when the path changes. This feature ensures that the device can update the member port information in IPv6 multicast forwarding entries and switch IPv6 multicast traffic to the new path promptly.

#### Examples

```
Enable the device to send MLD general queries upon a path change.
<Sysname> system-view
[Sysname] mld-snooping
[Sysname-mld-snooping] send-query enable
```

#### Related commands

**mld-snooping**

## New feature: Enabling embedded RP

### Restrictions and guidelines

This feature is supported only for IPv6 multicast group addresses in the range of FF7x::/12 or FFFx::/12 and in compliance with RFC 3956.

Enable this feature on all devices in the IPv6 PIM-SM domain.

IPv6 BIDIR-PIM does not support embedded RP.

## Procedure

1. Enter system view.  
**system-view**
2. Enter IPv6 PIM view.  
**ipv6 pim** [ **vpn-instance** *vpn-instance-name* ]
3. Enable embedded RP.  
**embedded-rp** [ *ipv6-acl-number* ]

By default, embedded RP is disabled.

## Command reference

Use **embedded-rp** to enable embedded RP.

Use **undo embedded-rp** to disable embedded RP.

### Syntax

```
embedded-rp [ipv6-acl-number]
undo embedded-rp
```

### Default

Embedded RP is disabled.

### Views

IPv6 PIM view

### Predefined user roles

network-admin

### Parameters

*ipv6-acl-number*: Specifies an IPv6 basic ACL number in the range of 2000 to 2999.

### Usage guidelines

Embedded RP is used for RP discovery in IPv6 PIM-SM. This feature supports embedding the RP address in an IPv6 multicast group address and enables multicast routers to extract the RP address from the IPv6 multicast group address.

Embedded RP is supported only for IPv6 multicast group addresses in the range of FF7x::/12 or FFFx::/12 and in compliance with RFC 3956.

If you do not specify the *ipv6-acl-number* argument, this command takes effect on all IPv6 multicast groups in the range of FF7x::/12 or FFFx::/12. If you specify this argument, this command takes effect only on the IPv6 multicast groups that are in the range of FF7x::/12 or FFFx::/12 and permitted by the ACL.

### Examples

# Enable embedded RP for IPv6 multicast group addresses in the range of FF7E:140:20::101/64.

```
<Sysname> system-view
[Sysname] acl ipv6 basic 2000
[Sysname-acl-ipv6-basic-2000] rule permit source ff7e:140:20::101 64
[Sysname-acl-ipv6-basic-2000] quit
[Sysname] ipv6 pim
```

## New feature: Specifying a local network to be advertised in the public instance or a VPN instance

### Specifying a local network to be advertised in the public instance or a VPN instance

#### About specifying a local network to be advertised

This feature specifies a local network of the public instance or a VPN instance. Then, you can configure BGP to redistribute the network (by using the **import-route local-aggregate** command) and advertise the network. For more information about the **import-route local-aggregate** command, see BGP commands in *Layer 3—IP Routing Command Reference*.

#### Restrictions and guidelines

The specified local network route must exist and be active in the routing table of the public instance or VPN instance.

#### Procedure

1. Enter system view.  
**system-view**
2. Enter public instance IPv4 address family view or VPN instance IPv4 address family view.
  - Execute the following commands in sequence to enter public instance IPv4 address family view:  
**ip public-instance**  
**address-family ipv4**
  - Execute the following commands in sequence to enter VPN instance IPv4 address family view:  
**ip vpn-instance** *vpn-instance-name*  
**address-family ipv4**
3. Specify a local network to be advertised.  
**network** *ipv4-address* [ *mask-length* | *mask* ]  
By default, no local network in the public instance or a VPN instance will be advertised.

## Command reference

### network

Use **network** to specify a local network to be advertised in the public instance or a VPN instance.

Use **undo network** to remove the local network to be advertised in the public instance or a VPN instance.

#### Syntax

**network** *ipv4-address* [ *mask-length* | *mask* ]

**undo network** *ipv4-address* [ *mask-length* | *mask* ]

## Default

No local network in the public instance or a VPN instance will be advertised.

## Views

Public instance IPv4 address family view

VPN instance IPv4 address family view

## Predefined user roles

network-admin

## Parameters

*ipv4-address*: Specifies a network by its IPv4 address. If you do not specify a mask or mask length, the natural network mask will be used.

*mask-length*: Specifies a mask length in the range of 0 to 32.

*mask*: Specifies a mask in dotted decimal notation.

## Usage guidelines

This command specifies a local network of the public instance or a VPN instance. Then, you can configure BGP to redistribute the network (by using the **import-route local-aggregate** command) and advertise the network.

The specified local network route must exist and be active in the routing table of the public instance or VPN instance.

## Examples

# Specify the local network to be advertised in VPN instance **vpn1** as network 10.0.0.0/16.

```
<Sysname> system-view
[Sysname] ip vpn-instance vpn1
[Sysname-vpn-instance-vpn1] address-family ipv4
[Sysname-vpn-ipv4-vpn1] network 10.0.0.0 255.255.0.0
```

## Related commands

**import-route** (*Layer 3—IP Routing Command Reference*)

# New feature: Changing the next hop address of VPNv4 routes to a VPN address

## About changing the next hop address of VPNv4 routes to a VPN address

By default, the device does not change the next hop of a received VPNv4 route. The next hop address of a VPNv4 route is a public address.

This feature can change the next hop address of VPNv4 routes received from the specified peer or peer group to an IP address in the corresponding VPN instances.

## Restrictions and guidelines

The device will delete and then re-establish the BGP session with the specified peer.

The device receives only the VPNv4 routes that have the same RDs as local existing RDs.

When the device sends a VPNv4 route, the device retains the original RT attribute of the route. It does not change the RT attribute of the route to that of the VPN instance.

If a VPN instance is deleted or the RD of a VPN instance is changed on the local device, the corresponding VPNv4 routes received by the device will be deleted.

## Procedure

1. Enter system view.  
**system-view**
2. Execute the following commands in sequence to enter BGP VPNv4 address family view:  
**bgp** *as-number* [ **instance** *instance-name* ]  
**address-family** **vpnvp4**
3. Change the next hop address of VPNv4 routes to a VPN address.  
**peer** { *group-name* | *ipv4-address* [ *mask-length* ] } **next-hop-vpn**  
By default, the device does not change the next hop attribute of a received VPNv4 route.

## Command reference

### peer next-hop-vpn

Use **peer next-hop-vpn** to change the next hop address of VPNv4 routes received from a peer or peer group to an IP address in the corresponding VPN instances.

Use **undo peer next-hop-vpn** to restore the default.

#### Syntax

```
peer { group-name | ipv4-address [mask-length] } next-hop-vpn
undo peer { group-name | ipv4-address [mask-length] } next-hop-vpn
```

#### Default

The device does not change the next hop attribute of a received VPNv4 route.

#### Views

BGP VPNv4 address family view

#### Predefined user roles

network-admin

#### Parameters

*group-name*: Specifies a peer group by its name, a case-sensitive string of 1 to 47 characters. The specified peer group must have existed.

*ipv4-address*: Specifies a peer by its IPv4 address. The specified peer must have existed.

*mask-length*: Specifies a mask length in the range of 0 to 32. You can use the *ip-address* and *mask-length* arguments together to specify a subnet. If you specify a subnet in this command, the device change the next hop of routes received from the dynamic peers in the subnet.

#### Usage guidelines

By default, the device does not change the next hop of a received VPNv4 route. The next hop address of a VPNv4 route is a public address.

This command changes the next hop address of VPNv4 routes received from the specified peer or peer group to an IP address in the corresponding VPN instances. The outgoing label of the VPNv4 route is also changed to an invalid value.

After this command is executed, the following applies:

- The device deletes and then re-establishes the BGP session with the specified peer.
- The device receives only the VPNv4 routes that have the same RDs as local existing RDs.

- When the device sends a VPNv4 route, the device retains the original RT attribute of the route. It does not change the RT attribute of the route to that of the VPN instance.
- If a VPN instance is deleted or the RD of a VPN instance is changed on the local device, the corresponding VPNv4 routes received by the device will be deleted.

## Examples

# In BGP VPNv4 address family view, configure the device to change the next hop of VPNv4 routes received from peer 1.1.1.1 to an IP address in the VPN instances.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] address-family vpnv4
[Sysname-bgp-default-vpnv4] peer 1.1.1.1 next-hop-vpn
```

# New feature: Displaying tunnel policy information

## Displaying tunnel policy information

Execute **display** commands in any view.

| Task                               | Command                                                    |
|------------------------------------|------------------------------------------------------------|
| Display tunnel policy information. | <b>display tunnel-policy</b> [ <i>tunnel-policy-name</i> ] |

## Command reference

### display tunnel-policy

Use **display tunnel-policy** to display tunnel policy information.

#### Syntax

```
display tunnel-policy [tunnel-policy-name]
```

#### Views

Any view

#### Predefined user roles

network-admin  
network-operator

#### Parameters

*tunnel-policy-name*: Specifies a tunnel policy by its name, a case-sensitive string of 1 to 19 characters. If you do not specify a tunnel policy, this command displays information about all tunnel policies.

## Examples

# Display information about all tunnel policies.

```
<Sysname> display tunnel-policy
Tunnel policy name : policy1
Preferred paths :
Interface : Not Configured
Tunnel type : Not Configured
```

```

Select-Seq : CRLSP
Load balance number : 1

Tunnel policy name : policy2
Preferred paths :
 Interface : Not Configured
 Tunnel type : Not Configured
Select-Seq : LSP
Load balance number : 8

Tunnel policy name : policy3
Preferred paths :
 Interface : Tunnel2
 Tunnel type : TE
Select-Seq : LSP CRLSP
Load balance number : 1

```

#### # Display information about tunnel policy **policy1**.

```

<Sysname> display tunnel-policy policy1
Tunnel policy name : policy1
Preferred paths :
 Interface : Not Configured
 Tunnel type : Not Configured
Select-Seq : CRLSP
Load balance number : 1

```

**Table 43 Command output**

| Field               | Description                                                                                                                                                  |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Preferred paths     | Information about preferred tunnels.                                                                                                                         |
| Interface           | Tunnel interface name of the preferred tunnel.                                                                                                               |
| Tunnel type         | Type of the preferred tunnel, which can be TE.                                                                                                               |
| Select-Seq          | Tunnel selection order. The tunnel types are displayed in descending order of priority in tunnel selection.<br>Supported tunnel types include CRLSP and LSP. |
| Load balance number | Number of tunnels for load balancing.                                                                                                                        |

## New feature: SBFD for MPLS

### Configuring SBFD for MPLS

#### About automatic SBFD for MPLS

SBFD for MPLS uses SBFD sessions to proactively verify the connectivity of LSP and MPLS TE tunnels. When a tunnel fails, SBFD can quickly detect the failure and notify the device to take an action, such as switching traffic to the backup tunnel.

SBFD is unidirectional and quicker than BFD. SBFD verifies the connectivity of one tunnel from the local device to the remote device.

SBFD for MPLS performs the following operations:

1. Establishes an SBFD session between the ingress and egress nodes of the tunnel to be inspected.
2. Adds the label associated with the tunnel into an SBFD control packet at the ingress (the initiator).
3. Sends the control packet to the egress node (the reflector) over the tunnel.
4. Determines the tunnel status according to whether the ingress node receives the SBFD control response from the egress node.

An SBFD session can be established only in static mode. You must specify the remote discriminator values for the SBFD sessions.

## Restrictions and guidelines

Perform this task on the SBFD session initiator. On the reflector, you must configure the **sbfd local-discriminator** command. Make sure the remote discriminator value configured on the initiator is the same as a local discriminator value configured on the reflector. If the two ends do not have a matching discriminator value, the reflector does not send responses to the initiator. For more information about the **sbfd local-discriminator** command, see *High Availability Command Reference*.

## Prerequisites

The source address of the SBFD session is the MPLS LSR ID of the local device. Before configuring SBFD for an LSP tunnel, perform the following tasks:

1. Configure an MPLS LSR ID for the local device.
2. Make sure a route is available on the remote device to reach the MPLS LSR ID.

## Procedure

1. Enter system view.  
**system-view**
2. Enable BFD for MPLS.  
**mpls bfd enable**  
By default, BFD for MPLS is disabled.
3. Enable SBFD to verify LSP or MPLS TE tunnel connectivity.
  - Configure SBFD to verify LSP connectivity for an FEC.  
**mpls sbfd dest-addr mask-length [ nexthop nexthop-address ] remote remote-id [ template template-name ]**  
By default, SBFD is not configured to verify LSP connectivity for an FEC.
  - Enter the view of an MPLS TE tunnel interface, enable SBFD to verify MPLS TE tunnel connectivity.  
**interface tunnel number**  
**mpls sbfd remote remote-id [ template template-name ]**  
By default, SBFD is not enabled for an MPLS TE tunnel.

## Command reference

### display mpls sbfd

Use **display mpls sbfd** to display SBFD information for LSP tunnels or MPLS TE tunnels.

## Syntax

```
display mpls sbfd [ipv4 ipv4-address mask-length | te tunnel
tunnel-number]
```

## Views

Any view

## Predefined user roles

network-admin  
network-operator

## Parameters

**ipv4** *ipv4-address mask-length*: Specifies an FEC by a destination IP address and mask length. The mask length is in the range of 0 to 32.

**te tunnel** *tunnel-number*: Specifies an MPLS TE tunnel by the tunnel interface number.

## Usage guidelines

If you do not specify any parameters, this command displays SBFD information for all LSP tunnels and all MPLS TE tunnels.

## Examples

# Display SBFD information for the LSPs to destination 22.22.2.2/32.

```
<Sysname> display mpls sbfd ipv4 22.22.2.2 32
```

```
Total number of sessions: 1, 1 up, 0 down, 0 init
```

```
FEC Type: LSP
```

```
FEC Info:
```

```
Destination: 22.22.2.2
```

```
Mask Length: 32
```

```
NHLFE ID: 1025
```

```
Local Discr: 513
```

```
Remote Discr: 513
```

```
Source IP: 11.11.1.1
```

```
Destination IP: 127.0.0.1
```

```
Session State: Up
```

```
Template Name: -
```

# Display SBFD information for MPLS TE tunnel 1.

```
<Sysname> display mpls sbfd te tunnel 1
```

```
Total number of sessions: 1, 1 up, 0 down, 0 init
```

```
FEC Type: TE Tunnel
```

```
FEC Info:
```

```
Source : 100.1.1.1
```

```
Destination: 200.1.1.1
```

```
Tunnel ID : 1
```

```
LSP ID : 100
```

```
NHLFE ID: 1025
```

```
Local Discr: 513
```

```
Remote Discr: 513
```

```
Source IP: 11.11.1.1
```

```
Destination IP: 127.0.0.1
```

```
Session State: Up
```

```
Template Name: -
```

**Table 44 Command output**

| Field          | Description                                                                                                                                                                                                        |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Local Discr    | Local discriminator of the SBFD session.                                                                                                                                                                           |
| Remote Discr   | Remote discriminator of the SBFD session.                                                                                                                                                                          |
| Source IP      | Source IP address of the SBFD session, which is the MPLS LSR ID of the local PE.                                                                                                                                   |
| Destination IP | Destination IP address of the SBFD session, which is an address on the 127.0.0.0/8 subnet.                                                                                                                         |
| Session State  | SBFD session state: <ul style="list-style-type: none"> <li>• <b>Init</b>—SBFD session is in initialization state.</li> <li>• <b>Up</b>—SBFD session is up.</li> <li>• <b>Down</b>—SBFD session is down.</li> </ul> |
| Template Name  | Name of the SBFD session parameter template.                                                                                                                                                                       |

**Related commands**

**mpls sbfd** (for LSP)

**mpls sbfd** (for TE tunnel)

**mpls sbfd (for LSP)**

Use **mpls sbfd** to enable SBFD for LSPs associated with an FEC.

Use **undo mpls sbfd** to disable SBFD for LSPs associated with an FEC.

**Syntax**

```
mpls sbfd dest-addr mask-length [nexthop nexthop-address] remote
remote-id [template template-name]
```

```
undo mpls sbfd dest-addr mask-length [nexthop nexthop-address]
```

**Default**

SBFD is not used to verify LSP connectivity for an FEC.

**Views**

System view

**Predefined user roles**

network-admin

**Parameters**

***dest-addr mask-length***: Specifies an FEC by a destination IP address and mask length. The mask length is in the range of 0 to 32.

***nexthop nexthop-address***: Specifies the next hop address of the FEC. If you specify the FEC next hop, SBFD verifies the specified LSP. If you do not specify a next hop, BFD verifies all LSPs for the FEC.

***remote remote-id***: Specifies the remote discriminator value for the SBFD session, in the range of 1 to 4294967295.

***template template-name***: Specifies an SBFD session parameter template by its name, a case-sensitive string of 1 to 63 characters. If you do not specify a template, SBFD uses the multihop SBFD session parameters configured in system view.

## Usage guidelines

Execute this command on the SBFD session initiator. After you enable BFD for MPLS and execute this command, the initiator creates an SBFD session for each LSP associated with the specified FEC. When an LSP fails, SBFD can quickly detect the failure and notify the initiator to take an action, such as switching traffic to the backup LSP.

An SBFD session for LSP connectivity verification can be established only in static mode. You must specify the remote discriminator value for the SBFD session. The remote discriminator value must be the same as a local discriminator value configured on the reflector by using the **sbfd local-discriminator** command. If the two ends do not have a matching discriminator value, the reflector does not send responses to the initiator.

The source address of the SBFD session is the MPLS LSR ID of the initiator. Before configuring SBFD for the LSP tunnel, perform the following tasks:

1. Configure an MPLS LSR ID for the initiator.
2. Make sure a route is available on the reflector to reach the MPLS LSR ID.

## Examples

# Use SBFD to verify the connectivity of the LSPs to destination 22.22.2.2/32.

```
<Sysname> system-view
[Sysname] mpls bfd enable
[Sysname] mpls sbfd 22.22.2.2 32 remote 1
```

# Use SBFD to verify the connectivity of the LSP with destination 22.22.2.2/32 and next hop 12.0.0.2. Set the remote discriminator to 1 and use SBFD session parameter template **test**.

```
<Sysname> system-view
[Sysname] mpls bfd enable
[Sysname] mpls sbfd 22.22.2.2 32 nexthop 12.0.0.2 remote 1 template test
```

## Related commands

**display mpls sbfd**

**mpls bfd enable**

**sbfd local-discriminator** (*High Availability Command Reference*)

## mpls sbfd (for TE tunnel)

Use **mpls sbfd** to enable SBFD for an MPLS TE tunnel.

Use **undo mpls sbfd** to restore the default.

## Syntax

**mpls sbfd remote** *remote-id* [**template** *template-name*]

**undo mpls sbfd**

## Default

SBFD is not enabled for an MPLS TE tunnel.

## Views

Tunnel interface view

## Predefined user roles

network-admin

## Parameters

**remote** *remote-id*: Specifies the remote discriminator value for the SBFD session, in the range of 1 to 4294967295.

**template** *template-name*: Specifies an SBFD session parameter template by its name, a case-sensitive string of 1 to 63 characters. If you do not specify a template, SBFD uses the multihop SBFD session parameters configured in tunnel interface view.

## Usage guidelines

Execute this command on the SBFD session initiator. After you enable BFD for MPLS and execute this command, the initiator creates an SBFD session on the MPLS TE tunnel. When the MPLS TE tunnel fails, SBFD can quickly detect the failure and notify the initiator to take an action, such as switching traffic to the backup tunnel.

An SBFD session for LSP connectivity verification can be established only in static mode. You must specify the remote discriminator value for the SBFD session. The remote discriminator value must be the same as a local discriminator value configured on the reflector by using the **sbfd local-discriminator** command. If the two ends do not have a matching discriminator value, the reflector does not send responses to the initiator.

The source address of the SBFD session is the MPLS LSR ID of the initiator. Before configuring SBFD on an MPLS TE tunnel, perform the following tasks:

1. Configure an MPLS LSR ID for the initiator.
2. Make sure a route is available on the reflector to reach the MPLS LSR ID.

## Examples

# Use SBFD to verify the connectivity of MPLS TE tunnel 1. Set the remote discriminator to 1 and use SBFD session parameter template **test**.

```
<Sysname> system-view
[Sysname] mpls bfd enable
[Sysname] interface Tunnel 1
[Sysname-Tunnel1] mpls sbfd remote 1 template test
```

## Related commands

**display mpls sbfd**

**mpls bfd enable**

**sbfd local-discriminator** (*High Availability Command Reference*)

# New feature: MPLS SR

## About MPLS SR

Segment Routing (SR) is a source routing technology. The source node selects a path for the packets, and then encodes the path in the packet header as an ordered list of segments. Each segment is identified by the segment identifier (SID). The SR nodes along the path forward the packets based on the SIDs in the packets. Only the source node needs to maintain the path status.

SR can operate with MPLS. In an MPLS network, SR uses MPLS labels as SIDs to forward packets on an LSP.

## MPLS SR characteristics

MPLS SR has the following characteristics:

- MPLS SR forwards packets based on the existing MPLS infrastructure. No infrastructure modifications are needed to implement SR on an MPLS network. For more information about the MPLS infrastructure, see MPLS basics configuration in *MPLS Configuration Guide*.
- MPLS SR expands and optimizes existing IGPs and uses the IGPs to distribute labels.
- MPLS SR implements network features such as MPLS TE in a simpler way, eliminating issues such as heavy and complicated routing protocol deployment.

## Basic concepts

- **SR node**—A node enabled with the MPLS SR feature. The ingress node (source node) adds labels to packets. The transit nodes forward packets based on labels. The egress node removes labels and forwards packets to the destination networks. SR nodes form an SR domain.
- **Segment**—An instruction an SR node executes on the incoming packet.
- **SID**—Segment ID, which is MPLS label in MPLS SR.
- **Segment type**—The following types of segments are available:
  - **Prefix segment**—SIDs are assigned to nodes based on destination address prefix. The nodes create prefix-specific forwarding entries.
  - **Adjacency segment**—SIDs are assigned to nodes based on adjacency.
- **SRLSP**—Segment routing label switched path, an LSP along which SR uses MPLS labels as SIDs to forward packets.
- **SRGB**—Segment routing global block, a range of global labels dedicated for MPLS SR prefix SIDs. Different nodes can have different SRGBs. The minimum label value of an SRGB label range is referred to as the base value of the SRGB. Labels 16000 through 24000 are reserved for SRGBs.

## How MPLS SR works

MPLS SR involves the following steps:

1. Label allocation for all nodes and links along the packet forwarding paths.  
[Dynamic SID allocation](#)
2. [Label forwarding entry installation based on SIDs](#). All MPLS SR devices in the SR domain use the allocated labels to create label forwarding entries.
3. [SRLSP setup](#). You can manually configure SRLSPs, or use an IGP or a controller to dynamically create SRLSPs.

When the ingress node of an SRLSP receives a packet, it adds labels to the packet and forwards the packet to the egress node through the SRLSP. After receiving a packet from the SRLSP, the egress node removes the label and forwards the packet based on the routing table.

You can bind a higher layer application (for example, MPLS TE) to an SRLSP to forward traffic of the application through the SRLSP.

## Dynamic SID allocation

MPLS SR expands IGP protocols such as IS-IS to advertise SIDs in IGP protocol packets. MPLS SR provides the following methods for dynamically allocating and advertising SIDs:

- **Prefix SID**—After you configure an SID for the loopback address on each SR node, the SIDs uniquely identify the SR nodes. All SR nodes in the SR domain use an IGP to advertise their own prefix SIDs. After receiving advertised prefix SIDs, each SR node calculates the prefix SIDs to the advertisers.

The prefix SID advertisement can be one of the following types:

- **Absolute value advertisement**—Each SR node advertises the prefix SID absolute value and the SRGB.
- **Index value advertisement**—Each SR node advertises the prefix SID index and the SRGB.

Each SR node is assigned a globally unique index value for the node's prefix. The prefix SID an SR node allocates for a prefix equals the SRGB base of the SR node plus the index for that prefix.

#### NOTE:

The device supports only the index value advertisement in the current software release.

- **Adjacency SID**—Each SR node with adjacency SID allocation enabled allocates SIDs for the IGP adjacencies (that is, the links to its IGP neighbors), and uses the SIDs to identify the links.

## Label forwarding entry installation based on SIDs

### Label forwarding entry installation based on prefix SIDs

Label forwarding entries based on prefix SIDs can be static or dynamic.

- **Static label forwarding entries**—The device creates local label forwarding entries based on manually configured incoming labels, outgoing labels, and next hops.
- **Dynamic label forwarding entries**—The device uses the IGP to flood in the SR domain the local SRGB and the index of the prefix SID for the local loopback interface. Based on the flooded information, the other devices in the domain calculate their local label forwarding entries by using the following rules:
  - Incoming label = Local SRGB base value + Index
  - Outgoing label = SRGB base value of the next hop + Index

**Figure 8 Creating label forwarding entries based on prefix SIDs**

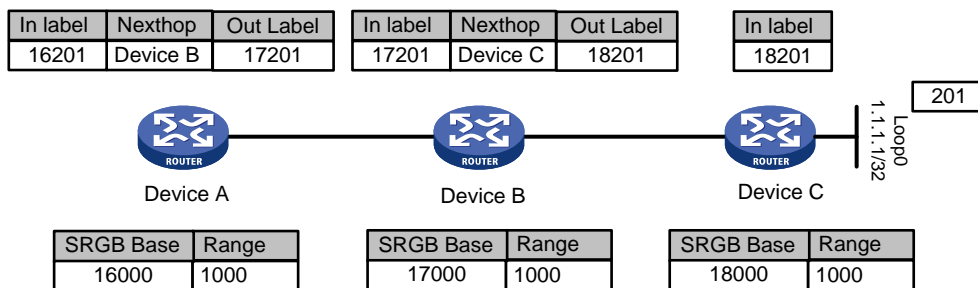


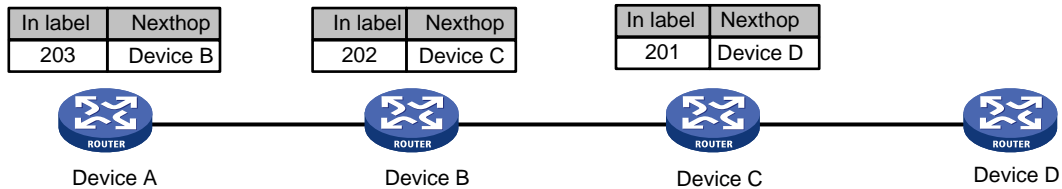
Figure 8 shows how dynamic label forwarding entries are created. After you assign index value 201 to loopback address 1.1.1.1/32 on Device C, Device C uses an IGP packet to advertise the index value and its local SRGB. Then, the devices calculate incoming and outgoing labels according to the previously mentioned label calculation rules.

- Device C calculates its incoming label, which is 18201.
- Device B calculates its incoming label and outgoing label and creates a label forwarding entry. The incoming label is 17201. The outgoing label is 18201. The next hop is Device C.
- Device A calculates its incoming label and outgoing label and creates a label forwarding entry. The incoming label is 16201. The outgoing label is 17201. The next hop is Device B.

### Label forwarding entry installation based on adjacency SIDs

When using adjacency SIDs, each device allocates a static or dynamic incoming label to the link to its neighbor. The label has local significance. Multiple devices can use the same adjacency SID.

**Figure 9 Creating label forwarding entries based on adjacency SIDs**



As shown in [Figure 9](#), the devices are running the same IGP. After IGP adjacencies are established between the devices, Device A, Device B, and Device C allocates labels and creates label forwarding entries as follows:

- Device A allocates label 203 to the link to its neighbor Device B.
- Device B allocates label 202 to the link to its neighbor Device C.
- Device C allocates label 201 to the link to its neighbor Device D.
- Device A creates a label forwarding entry with incoming label 203 and next hop Device B.
- Device B creates a label forwarding entry with incoming label 202 and next hop Device C.
- Device C creates a label forwarding entry with incoming label 201 and next hop Device D.

## SRLSP setup

You can use the following methods to create SRLSPs:

- Manually configure an SRLSP.  
To configure an SRLSP, you must specify the label stack for packets to be forwarded along the SRLSP on the ingress node. Each label in the stack corresponds to a prefix SID or adjacency SID. The system can look for the outgoing label and next hop based on the prefix SID or adjacency SID.
- Configure SR nodes to use an IGP to dynamically establish an SRLSP.  
SR nodes follow these steps to establish SRLSPs:
  - a. Use an IGP to collect prefix SID information from the MPLS SR network.
  - b. Calculate the shortest paths to other SR nodes based on the collected prefix SID information and the IGP network topology.
  - c. Establish SRLSPs based on the shortest paths.
- Configure a controller to deploy SRLSP configuration to the device so the device creates SRLSPs.  
For more information about controller configuration, see the user guide for the controller.

## Packet forwarding in MPLS SR

Based on the SID allocation method, MPLS SR uses one of the following packet forwarding methods:

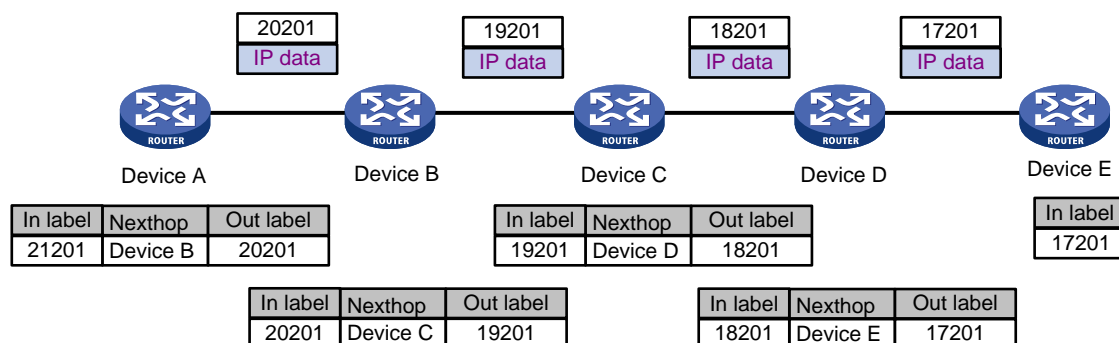
- **Prefix SID-based packet forwarding**—The ingress node encapsulates the prefix SID for the egress node to a packet. The transit nodes forward the packet based on label forwarding entries.
- **Adjacency SID-based packet forwarding**—The ingress node encapsulates the label stack that contains the adjacency SIDs of all links along the forwarding path to a packet. Each transit node uses the top label in the label stack to determine the next hop and pops the top label before forwarding the packet to the next hop.
- **Prefix and adjacency SID-based packet forwarding**—The nodes use prefix SID-based packet forwarding in combination with adjacency-based packet forwarding.

## Prefix SID-based packet forwarding

Figure 10 shows how a packet is forwarded along the SRLSP from Device A to Device E based on prefix SIDs. In this example, the outgoing label for the packet is 21201 on Device A.

1. Ingress node Device A searches for a forwarding entry for label 21201, adds outgoing label 20201 to the packet and sends the packet to the next hop (Device B).
2. When transit node Device B receives the packet, it searches for a label forwarding entry that matches the label in the packet. Then, Device B uses the outgoing label of the matched entry (19201) to replace the label in the packet and forwards the packet to the next hop (Device C).
3. Transit nodes Device C and Device D process the packet in the same way Device B does.
4. When egress node Device E receives the packet, it removes the label and forwards the packet by IP address.

**Figure 10 Prefix SID-based packet forwarding**

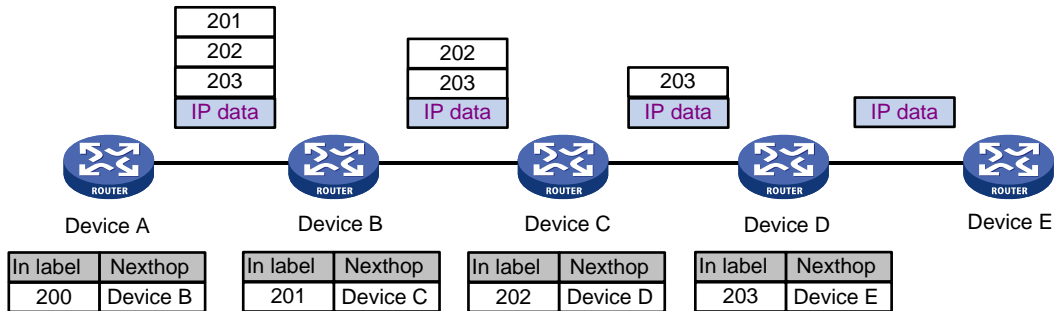


## Adjacency SID-based packet forwarding

Figure 11 shows how a packet is forwarded along the SRLSP from Device A to Device E based on adjacency SIDs. In this example, the label stack for the packet is configured as (200, 201, 202, and 203) on Device A.

1. Ingress node Device A searches for a forwarding entry for the top label (200) to determine the next hop. Then, Device A adds label stack (201, 202, and 203) to the packet, and forwards the packet to the next hop (Device B).
2. When transit node Device B receives the packet, it searches for a forwarding entry for the top label (201) to determine the next hop. Then, Device B removes the top label from the stack and forwards the packet to the next hop (Device C).
3. When transit node Device C receives the packet, it searches for a forwarding entry for the top label (202) to determine the next hop. Then, Device C removes the top label from the stack and forwards the packet to the next hop (Device D).
4. When transit node Device D receives the packet, it searches for a forwarding entry for the label (203) to determine the next hop. Then, Device D removes the label stack from the packet and forwards the packet to the next hop (Device E).
5. When egress node Device E receives the packet, it forwards the packet by IP address.

**Figure 11 Adjacency SID-based packet forwarding**

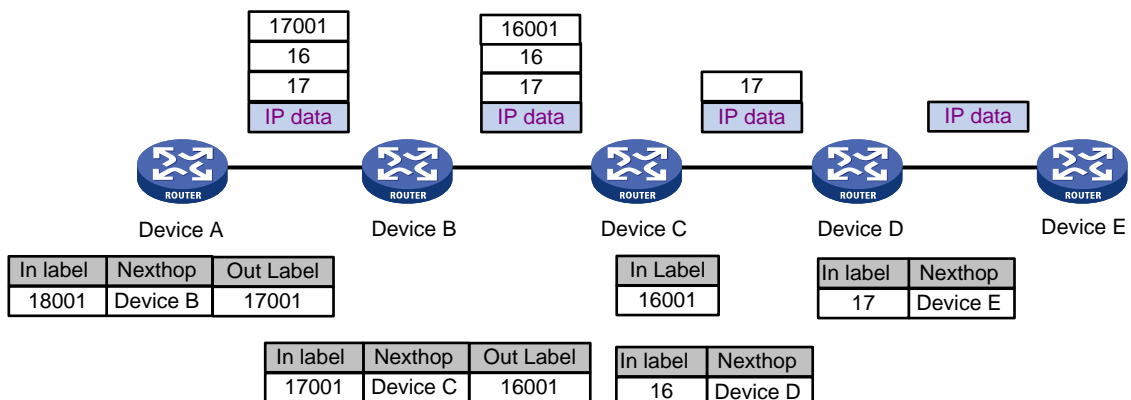


## Prefix and adjacency SID-based packet forwarding

Figure 12 shows how a packet is forwarded along the SRLSP from Device A to Device E based on prefix SIDs and adjacency SIDs. In this example, the index value for the prefix SID of Device C is 1. The prefix SIDs for Device A, Device B, and Device C are 18001, 17001, and 16001, respectively. The Adjacency SIDs that Device C and Device D allocate for the adjacent links are 16 and 17, respectively. On Device A, the label stack for the packet is (18001, 16, 17).

1. Ingress node Device A searches for a forwarding entry for label 18001 to determine the outgoing label (17001) and next hop (Device B). Device A swaps label 18001 with 17001. Then, it adds label stack (17001, 16, 17) to the packet and sends the packet to the next hop (Device B).
2. When transit node Device B receives the packet, it searches for a label forwarding entry that matches the top label in the label stack (17001). Then, Device B uses the outgoing label of the matched entry (16001) to replace the top label and forwards the packet to the next hop (Device C).
3. When transit node Device C receives the packet, it removes the top label 16001 and searches for a forwarding entry for the next label (16) to determine the next hop. Then, Device C removes label 16 from the stack and forwards the packet to the next hop (Device D).
4. When transit node Device D receives the packet, it searches for a forwarding entry for the label (17) to determine the next hop. Then, Device D removes the label stack from the packet and forwards the packet to the next hop (Device E).
5. When egress node Device E receives the packet, it forwards the packet by IP address.

**Figure 12 Prefix and adjacency SID-based packet forwarding**



## MPLS SR and LDP interworking

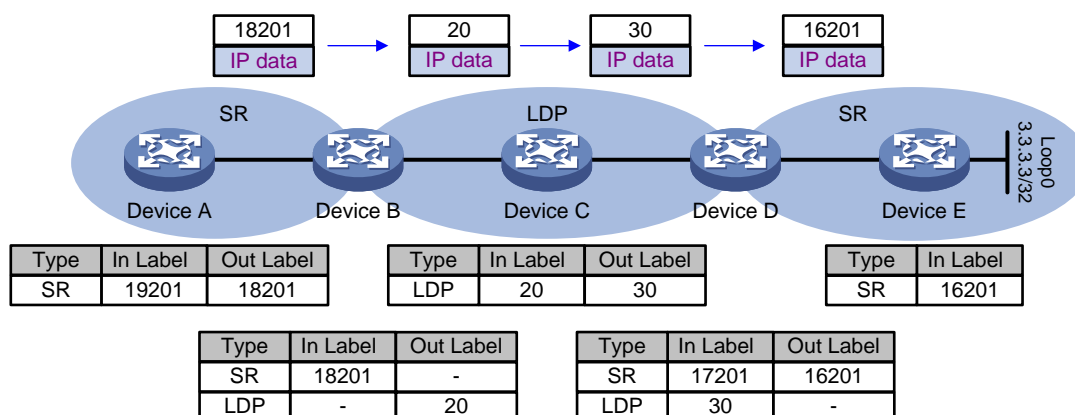
### ! IMPORTANT:

For an MPLS SR network to communicate with an LDP network, make sure SRLSPs use prefix SIDs.

For an MPLS SR network and an LDP network to communicate, the SR/LDP border devices must be able to connect the SR LSP and the LDP LSP as follows:

- **MPLS SR to LDP interworking**—The border device installs SR-to-LDP label forwarding entries. For a packet from the MPLS SR network to the LDP network, the MPLS SR label forwarding entry on the border device does not have an outgoing label. The border device uses the outgoing label of the LDP label forwarding entry for the same destination address as the outgoing label of the packet.
- **LDP to MPLS SR interworking**—The border device installs LDP-to-SR label forwarding entries. For a packet from the LDP network to the MPLS SR network, the LDP forwarding entry on the border device does not have an outgoing label. The border device must use the outgoing label of the MPLS SR forwarding entry for the same destination address as the outgoing label of the packet.

Figure 13 MPLS SR and LDP interworking



As shown in Figure 13, Device A, Device B, Device D, and Device E are running MPLS SR. After you assign index value 201 to loopback interface address 3.3.3.3/32 on Device E, Device E will advertise the index value and the local SRGB. After Device A, Device B, and Device D receive the message, they will install their respective MPLS SR label forwarding entries. Device B, Device C, and Device D are running LDP. They allocate labels to destination address 3.3.3.3/32 to generate the LDP label forwarding entries.

A packet that Device A sends to Device E will be forwarded as follows:

2. Ingress node Device A encapsulates label 18201 to the packet and forwards the packet to transit node Device B.
3. Transit node Device B searches for an MPLS SR label forwarding entry for incoming label 18201 and finds that the entry does not have an outgoing label. Because an LDP label forwarding entry with outgoing label 20 exists for the destination address (3.3.3.3/32), Device B encapsulates outgoing label 20 in the packet. Then Device B sends the packet to transit node Device C.
4. Device C forwards the packet to Device D based on its LDP label forwarding entries. The outgoing label is 30.
5. Device D searches for an LDP label forwarding entry for incoming label 30 and finds that the entry does not have an outgoing label. Because an MPLS SR label forwarding entry with outgoing label 16201 exists for the destination address (3.3.3.3/32), Device B encapsulates outgoing label 16201 in the packet. Then Device D sends the packet to egress node Device E.

6. Egress node Device E deletes the incoming label and forwards the packet by IP address.

## Protocols and standards

- *draft-ietf-spring-segment-routing-mpls-00*
- *draft-ietf-spring-segment-routing-02*
- *draft-ietf-isis-segment-routing-extensions-06*
- *draft-ietf-spring-segment-routing-11*

## MPLS SR tasks at a glance

### IP traffic forwarding over SRLSPs tasks at a glance

To forward IP traffic over SRLSPs, perform the following configuration tasks:

1. Configuring segments
  - [Configuring IGP-based SID advertisement](#)
2. (Optional.) [Configuring the device to prefer SRLSPs in traffic forwarding](#)

## Configuring IGP-based SID advertisement

### IGP-based SID advertisement tasks at a glance

Perform the following tasks on each node along an SRLSP:

1. [Configuring the IGP to support MPLS SR](#)
2. Perform one or more of the following tasks:
  - [Configuring prefix SIDs](#)
  - [Enabling MPLS SR adjacency label allocation for the IGP](#)
3. [Configuring the MPLS SRGB](#)  
Perform this task if you are configuring prefix SIDs.

## Prerequisites

Before you configure IGP-based SID advertisement, perform the following tasks:

- Determine the ingress node, transit nodes, and egress node of an SRLSP.
- Determine the following items for each node:
  - Index value for the prefix SID.
  - MPLS SRGB.
- Enable MPLS on all nodes and interfaces that will participate in MPLS forwarding. For information about enabling MPLS, see "Configuring basic MPLS."

## Configuring the IGP to support MPLS SR

### Prerequisites

For MPLS SR to take effect, set the IS-IS cost style to wide, compatible, or wide-compatible before configuring IS-IS to support MPLS SR. For more information about the cost style, see IS-IS configuration in *Layer 3—IP Routing Configuration Guide*.

## Configuring IS-IS to support MPLS SR

1. Enter system view.  
**system-view**
  2. Enter IS-IS view.  
**isis** *process-id*
  3. Enter IS-IS IPv4 unicast address family view.  
**address-family** **ipv4**
  4. Enable MPLS SR.  
**segment-routing** **mpls**
- By default, MPLS SR is disabled.

## Configuring prefix SIDs

### About prefix SID configuration

Configuring a prefix SID in loopback interface view binds the SID with the IP address of the loopback interface.

To configure a prefix SID, use one of the following methods:

- Specify an absolute value in the SRGB. The absolute value is used as the prefix SID.
- Specify an index value. The sum of the index value and the SRGB base value is used as the prefix SID. The prefix SID must be in the SRGB.

### Restrictions and guidelines

To use a prefix SID for a group of SR nodes in anycast scenarios, specify the **n-flag-clear** keyword to set the Node-SID flag bit of the prefix SID to 0.

To configure an IS-IS prefix SID, you must enable an IS-IS process on the loopback interface.

### Configuring an IS-IS prefix SID

1. Enter system view.  
**system-view**
  2. Enter loopback interface view.  
**interface** **loopback** *interface-number*
  3. Configure an IS-IS prefix SID.  
**isis** **prefix-sid** { **absolute** *absolute-value* | **index** *index-value* }  
[ **n-flag-clear** ] [ **explicit-null** ]
- By default, no IS-IS prefix SID is configured.

## Enabling MPLS SR adjacency label allocation for the IGP

### Restrictions and guidelines

For this feature to take effect, you must enable MPLS SR.

### Enabling MPLS SR adjacency label allocation for IS-IS

1. Enter system view.  
**system-view**
2. Enter IS-IS view.  
**isis** *process-id*
3. Enter IS-IS IPv4 unicast address family view.

**address-family ipv4**

4. Enable MPLS SR adjacency label allocation.

**segment-routing adjacency enable**

By default, MPLS SR adjacency label allocation is disabled.

## Configuring the MPLS SRGB

### Restrictions and guidelines

To configure the SRGB of a node successfully, make sure the SRGB contains the prefix SID configured for the node.

MPLS reserves labels in the range of 16000 to 24000 for SRGB. The minimum and maximum label values of the SRBG must be in the range.

### Configuring the MPLS SRGB for IS-IS

1. Enter system view.

**system-view**

2. Enter IS-IS view.

**isis** *process-id*

3. Configure the MPLS SRGB.

**segment-routing global-block** *minvalue maxvalue*

By default, the minimum label value is 16000, and the maximum label value is 24000.

## Configuring the device to prefer SRLSPs in traffic forwarding

### About preferring SRLSPs in traffic forwarding

This feature enables the device to preferentially use SRLSPs to forward traffic when both SRLSPs and LDP LSPs are available for traffic forwarding. If you do not configure this feature, the device prefers to use LDP LSPs for traffic forwarding.

### Restrictions and guidelines

This feature takes effect only when MPLS SR is enabled and SRLSPs use prefix SIDs.

### Configuring the device to prefer SRLSPs established by IS-IS in traffic forwarding

1. Enter system view.

**system-view**

2. Enter IS-IS view.

**isis** *process-id*

3. Configure the device to prefer SRLSPs in traffic forwarding.

**segment-routing sr-prefer** [ **prefix-list** *prefix-list-name* ]

By default, the device prefers LDP LSPs to SRLSPs.

## Display and maintenance commands for MPLS SR

Execute **display** commands in any view.

| Task                                            | Command                                                                |
|-------------------------------------------------|------------------------------------------------------------------------|
| Display IS-IS SR adjacency segment information. | <b>display isis segment-routing adjacency</b><br>[ <i>process-id</i> ] |

| Task                            | Command                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------|
| Display IS-IS SRGB information. | <b>display isis segment-routing global-block</b><br>[ <b>level-1</b>   <b>level-2</b> ] [ <i>process-id</i> ] |

## MPLS SR configuration examples

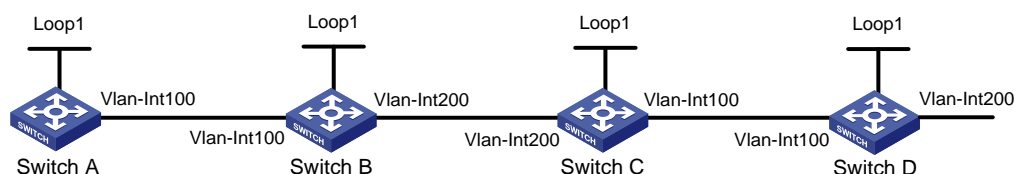
### Example: Configuring MPLS SR based on ISIS-advertised SIDs

#### Network configuration

As shown in Figure 14, Switch A, Switch B, Switch C and Switch D are running IS-IS.

Configure dynamic SID allocation on loopback interfaces of the switches. Then, establish an SRLSP from Switch A to Switch D based on the allocated SIDs and configure an MPLS TE tunnel over the SRLSP to transmit data.

**Figure 14 Network diagram**



**Table 45 Interface and IP address assignment**

| Device   | Interface   | IP address  | Device   | Interface   | IP address   |
|----------|-------------|-------------|----------|-------------|--------------|
| Switch A | Loop1       | 1.1.1.1/32  | Switch B | Loop1       | 2.2.2.2/32   |
|          | Vlan-Int100 | 10.0.0.1/24 |          | Vlan-Int100 | 10.0.0.2/24  |
|          |             |             |          | Vlan-Int200 | 11.0.0.1/24  |
| Switch C | Loop1       | 3.3.3.3/32  | Switch D | Loop1       | 4.4.4.4/32   |
|          | Vlan-Int100 | 12.0.0.1/24 |          | Vlan-Int100 | 12.0.0.2/24  |
|          | Vlan-Int200 | 11.0.0.2/24 |          | Vlan-Int200 | 100.1.2.1/24 |

#### Procedure

1. Configure IP addresses and masks for interfaces. (Details not shown.)
2. Configure IS-IS on the switches and set the IS-IS cost style to wide:

# Configure Switch A.

```
<SwitchA> system-view
[SwitchA] isis 1
[SwitchA-isis-1] network-entity 00.0000.0000.0001.00
[SwitchA-isis-1] cost-style wide
[SwitchA-isis-1] quit
[SwitchA] interface vlan-interface 100
[SwitchA-vlan-interface100] isis enable 1
[SwitchA-vlan-interface100] quit
[SwitchA] interface loopback 1
[SwitchA-LoopBack1] isis enable 1
[SwitchA-LoopBack1] quit
```

# Configure Switch B.

```

<SwitchB> system-view
[SwitchB] isis 1
[SwitchB-isis-1] network-entity 00.0000.0000.0002.00
[SwitchB-isis-1] cost-style wide
[SwitchB-isis-1] quit
[SwitchB] interface vlan-interface 100
[SwitchB-vlan-interface100] isis enable 1
[SwitchB-vlan-interface100] quit
[SwitchB] interface vlan-interface 200
[SwitchB-vlan-interface200] isis enable 1
[SwitchB-vlan-interface200] quit
[SwitchB] interface loopback 1
[SwitchB-LoopBack1] isis enable 1
[SwitchB-LoopBack1] quit

```

#### # Configure Switch C.

```

<SwitchC> system-view
[SwitchC] isis 1
[SwitchC-isis-1] network-entity 00.0000.0000.0003.00
[SwitchC-isis-1] cost-style wide
[SwitchC-isis-1] quit
[SwitchC] interface vlan-interface 100
[SwitchC-vlan-interface100] isis enable 1
[SwitchC-vlan-interface100] quit
[SwitchC] interface vlan-interface 200
[SwitchC-vlan-interface200] isis enable 1
[SwitchC-vlan-interface200] quit
[SwitchC] interface loopback 1
[SwitchC-LoopBack1] isis enable 1
[SwitchC-LoopBack1] quit

```

#### # Configure Switch D.

```

<SwitchD> system-view
[SwitchD] isis 1
[SwitchD-isis-1] network-entity 00.0000.0000.0004.00
[SwitchD-isis-1] cost-style wide
[SwitchD-isis-1] quit
[SwitchD] interface vlan-interface 100
[SwitchD-vlan-interface100] isis enable 1
[SwitchD-vlan-interface100] quit
[SwitchD] interface vlan-interface 200
[SwitchD-vlan-interface200] isis enable 1
[SwitchD-vlan-interface200] quit
[SwitchD] interface loopback 1
[SwitchD-LoopBack1] isis enable 1
[SwitchD-LoopBack1] quit

```

### 3. Configure LSR IDs, and enable MPLS and MPLS TE:

#### # Configure Switch A.

```

[SwitchA] mpls lsr-id 1.1.1.1
[SwitchA] mpls te

```

```
[SwitchA-te] quit
[SwitchA] interface vlan-interface 100
[SwitchA-vlan-interface100] mpls enable
[SwitchA-vlan-interface100] quit
```

#### # Configure Switch B.

```
[SwitchB] mpls lsr-id 2.2.2.2
[SwitchB] mpls te
[SwitchB-te] quit
[SwitchB] interface vlan-interface 100
[SwitchB-vlan-interface100] mpls enable
[SwitchB-vlan-interface100] quit
[SwitchB] interface vlan-interface 200
[SwitchB-vlan-interface200] mpls enable
[SwitchB-vlan-interface200] quit
```

#### # Configure Switch C.

```
[SwitchC] mpls lsr-id 3.3.3.3
[SwitchC] mpls te
[SwitchC-te] quit
[SwitchC] interface vlan-interface 100
[SwitchC-vlan-interface100] mpls enable
[SwitchC-vlan-interface100] quit
[SwitchC] interface vlan-interface 200
[SwitchC-vlan-interface200] mpls enable
[SwitchC-vlan-interface200] quit
```

#### # Configure Switch D.

```
[SwitchD] mpls lsr-id 4.4.4.4
[SwitchD] mpls te
[SwitchD-te] quit
[SwitchD] interface vlan-interface 100
[SwitchD-vlan-interface100] mpls enable
[SwitchD-vlan-interface100] quit
```

### 4. Configure SRGBs and enable MPLS SR on the switches:

#### # Configure Switch A.

```
[SwitchA] isis 1
[SwitchA-isis-1] segment-routing global-block 16000 16999
[SwitchA-isis-1] address-family ipv4
[SwitchA-isis-1-ipv4] segment-routing mpls
[SwitchA-isis-1-ipv4] quit
[SwitchA-isis-1] quit
```

#### # Configure Switch B.

```
[SwitchB] isis 1
[SwitchB-isis-1] segment-routing global-block 17000 17999
[SwitchB-isis-1] address-family ipv4
[SwitchB-isis-1-ipv4] segment-routing mpls
[SwitchB-isis-1-ipv4] quit
[SwitchB-isis-1] quit
```

#### # Configure Switch C.

```
[SwitchC] isis 1
```

```
[SwitchC-isis-1] segment-routing global-block 18000 18999
[SwitchC-isis-1] address-family ipv4
[SwitchC-isis-1-ipv4] segment-routing mpls
[SwitchC-isis-1-ipv4] quit
[SwitchC-isis-1] quit
```

#### # Configure Switch D.

```
[SwitchD] isis 1
[SwitchD-isis-1] segment-routing global-block 19000 19999
[SwitchD-isis-1] address-family ipv4
[SwitchD-isis-1-ipv4] segment-routing mpls
[SwitchD-isis-1-ipv4] quit
[SwitchD-isis-1] quit
```

### 5. Configure IS-IS prefix SIDs for the switches:

#### # Configure Switch A.

```
[SwitchA] interface loopback 1
[SwitchA-LoopBack1] isis prefix-sid index 10
```

#### # Configure Switch B.

```
[SwitchB] interface loopback 1
[SwitchB-LoopBack1] isis prefix-sid index 20
```

#### # Configure Switch C.

```
[SwitchC] interface loopback 1
[SwitchC-LoopBack1] isis prefix-sid index 30
```

#### # Configure Switch D.

```
[SwitchD] interface loopback 1
[SwitchD-LoopBack1] isis prefix-sid index 40
```

### 6. On Switch A, establish static SRLSP static-sr-lsp-1 to Switch D and configure an MPLS TE tunnel over the static SRLSP:

# Configure Switch A as the ingress node of static SRLSP **static-sr-lsp-1** and specify the prefix label that Switch A allocated for Switch D (16040) as the outgoing label.

```
[SwitchA] static-sr-mpls lsp static-sr-lsp-1 out-label 16040
```

# Establish static MPLS TE tunnel 1 to Switch D and specify the LSR ID of Switch D as the tunnel destination address. Bind static SRLSP **static-sr-lsp-1** to MPLS TE tunnel interface 1.

```
[SwitchA] interface tunnel 1 mode mpls-te
[SwitchA-Tunnel1] ip address 6.1.1.1 255.255.255.0
[SwitchA-Tunnel1] destination 4.4.4.4
[SwitchA-Tunnel1] mpls te signaling static
[SwitchA-Tunnel1] mpls te static-sr-mpls static-sr-lsp-1
[SwitchA-Tunnel1] quit
```

### 7. On Switch A, configure a static route to direct traffic destined for 100.1.2.0/24 to MPLS TE tunnel 1.

```
[SwitchA] ip route-static 100.1.2.0 24 tunnel 1 preference 1
```

## Verifying the configuration

#### # Display IS-IS process information on Switch A.

```
[SwitchA] display isis
```

```
IS-IS(1) Protocol Information
```

```

Network entity : 00.0000.0000.0001.00
IS level : level-1
Cost style : Wide
Fast reroute : Disabled
Preference : 15
LSP length receive : 1497
LSP length originate
 level-1 : 1497
Maximum imported routes : 100000
Timers
 LSP-max-age : 1200
 LSP-refresh : 900
 SPF mode : Normal
 SPF intervals : 5 50 200

```

#### Segment routing

```

MPLS : Enabled
Adjacency : Disabled
global block : 16000 16999

```

# Display detailed IS-IS interface information on Switch A to view SID information for the loopback interface.

```
[SwitchA] display isis interface verbose
```

#### Interface information for IS-IS(1)

```
Interface: LoopBack1
```

| Index                       | IPv4 state | IPv6 state       | Circuit ID  | MTU  | Type  | DIS |
|-----------------------------|------------|------------------|-------------|------|-------|-----|
| 00002                       | Up         | Down             | 1           | 1536 | L1/L2 | --  |
| SNPA address                |            | : 0000-0000-0000 |             |      |       |     |
| IP address                  |            | : 1.1.1.1        |             |      |       |     |
| Secondary IP address(es)    |            | :                |             |      |       |     |
| IPv6 link-local address     |            | :                |             |      |       |     |
| Extended circuit ID         |            | : 2              |             |      |       |     |
| CSNP timer value            |            | : L1             | 10          | L2   | 10    |     |
| Hello timer value           |            | :                | 10          |      |       |     |
| Hello multiplier value      |            | :                | 3           |      |       |     |
| LSP timer value             |            | : L12            | 33          |      |       |     |
| LSP transmit-throttle count |            | : L12            | 5           |      |       |     |
| Cost                        |            | : L1             | 0           | L2   | 0     |     |
| IPv6 cost                   |            | : L1             | 0           | L2   | 0     |     |
| Priority                    |            | : L1             | 64          | L2   | 64    |     |
| Retransmit timer value      |            | : L12            | 5           |      |       |     |
| MPLS TE status              |            | : L1 Disabled    | L2 Disabled |      |       |     |
| IPv4 BFD                    |            | : Disabled       |             |      |       |     |
| IPv6 BFD                    |            | : Disabled       |             |      |       |     |
| IPv4 FRR LFA backup         |            | : Enabled        |             |      |       |     |
| IPv6 FRR LFA backup         |            | : Enabled        |             |      |       |     |
| IPv4 prefix suppression     |            | : Disabled       |             |      |       |     |
| IPv6 prefix suppression     |            | : Disabled       |             |      |       |     |

```

IPv4 tag : 0
IPv6 tag : 0
Prefix-SID type : Index
Value : 10
Prefix-SID validity : Valid

```

#### # Display SRGB information on Switch A.

```
[SwitchA] display isis segment-routing global-block
```

```
Segment routing global block information for IS-IS(1)
```

```

```

```
Level-1 SRGB
```

```

```

| System ID      | Base  | Range |
|----------------|-------|-------|
| 0000.0000.0001 | 16000 | 1000  |
| 0000.0000.0002 | 17000 | 1000  |
| 0000.0000.0003 | 18000 | 1000  |
| 0000.0000.0004 | 19000 | 1000  |

#### # Display detailed IS-IS routing information on Switch A to view information about routes bound with labels.

```
[SwitchA] display isis route verbose
```

```
Route information for IS-IS(1)
```

```

```

```
Level-1 IPv4 Forwarding Table
```

```

```

|                         |                         |                  |
|-------------------------|-------------------------|------------------|
| IPv4 Dest : 10.0.0.0/24 | Int. Cost : 10          | Ext. Cost : NULL |
| Admin Tag : -           | Src Count : 2           | Flag : D/L/-     |
| InLabel : 4294967295    | InLabel Flag: -/-/-/-/- |                  |
| NextHop :               | Interface :             | ExitIndex :      |
| Direct                  | Vlan100                 | 0x00000102       |
| Nib ID : 0x0            | OutLabel : 4294967295   | OutLabelFlag: -  |
|                         |                         |                  |
| IPv4 Dest : 11.0.0.0/24 | Int. Cost : 20          | Ext. Cost : NULL |
| Admin Tag : -           | Src Count : 2           | Flag : R/-/-     |
| InLabel : 4294967295    | InLabel Flag: -/-/-/-/- |                  |
| NextHop :               | Interface :             | ExitIndex :      |
| 10.0.0.2                | Vlan100                 | 0x00000102       |
| Nib ID : 0x14000004     | OutLabel : 4294967295   | OutLabelFlag: -  |
|                         |                         |                  |
| IPv4 Dest : 12.0.0.0/24 | Int. Cost : 30          | Ext. Cost : NULL |
| Admin Tag : -           | Src Count : 2           | Flag : R/-/-     |
| InLabel : 4294967295    | InLabel Flag: -/-/-/-/- |                  |
| NextHop :               | Interface :             | ExitIndex :      |

```

10.0.0.2 Vlan100 0x00000102
Nib ID : 0x14000004 OutLabel : 4294967295 OutLabelFlag: -

IPv4 Dest : 1.1.1.1/32 Int. Cost : 0 Ext. Cost : NULL
Admin Tag : - Src Count : 1 Flag : D/L/-
InLabel : 16010 InLabel Flag: -/N/-/-/-/-
NextHop : Interface : ExitIndex :
Direct : Loop1 : 0x00000584
Nib ID : 0x0 OutLabel : 4294967295 OutLabelFlag: -

IPv4 Dest : 2.2.2.2/32 Int. Cost : 10 Ext. Cost : NULL
Admin Tag : - Src Count : 1 Flag : R/-/-
InLabel : 16020 InLabel Flag: -/N/-/-/-/-
NextHop : Interface : ExitIndex :
10.0.0.2 Vlan100 0x00000102
Nib ID : 0x14000003 OutLabel : 17020 OutLabelFlag: I

IPv4 Dest : 3.3.3.3/32 Int. Cost : 20 Ext. Cost : NULL
Admin Tag : - Src Count : 1 Flag : R/-/-
InLabel : 16030 InLabel Flag: -/N/-/-/-/-
NextHop : Interface : ExitIndex :
10.0.0.2 Vlan100 0x00000102
Nib ID : 0x14000002 OutLabel : 17030 OutLabelFlag: -

IPv4 Dest : 4.4.4.4/32 Int. Cost : 20 Ext. Cost : NULL
Admin Tag : - Src Count : 1 Flag : R/-/-
InLabel : 16040 InLabel Flag: -/N/-/-/-/-
NextHop : Interface : ExitIndex :
10.0.0.2 Vlan100 0x00000102
Nib ID : 0x14000002 OutLabel : 17040 OutLabelFlag: -

```

Flags: D-Direct, R-Added to Rib, L-Advertised in LSPs, U-Up/Down Bit Set

InLabel flags: R-Readvertisement, N-Node SID, P-no PHP  
E-Explicit null, V-Value, L-Local

OutLabelFlags: E-Explicit null, I-Implicit null, N-Normal

#### # Display MPLS LSP information on Switch A.

```
[SwitchA] display mpls lsp
```

| FEC        | Proto | In/Out Label | Interface/Out NHLFE |
|------------|-------|--------------|---------------------|
| 10.0.0.2   | Local | -/-          | Vlan100             |
| 1.1.1.1/32 | ISIS  | 16010/-      | -                   |
| 2.2.2.2/32 | ISIS  | 16020/3      | Vlan100             |
| 2.2.2.2/32 | ISIS  | -/3          | Vlan100             |
| 3.3.3.3/32 | ISIS  | 16030/17030  | Vlan100             |
| 3.3.3.3/32 | ISIS  | -/17030      | Vlan100             |
| 4.4.4.4/32 | ISIS  | 16040/17040  | Vlan100             |
| 4.4.4.4/32 | ISIS  | -/17040      | Vlan100             |

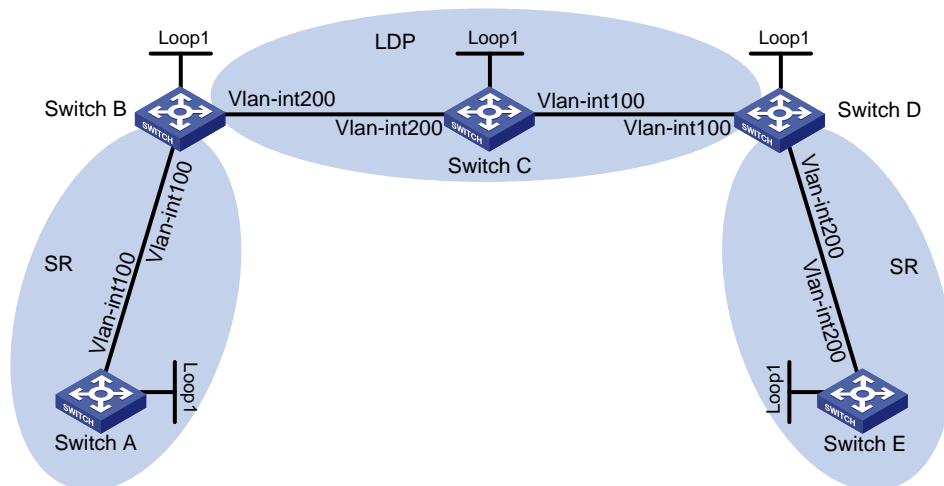
## Example: Configuring MPLS SR and LDP interworking

### Network configuration

As shown in [Figure 15](#), complete the following tasks so the two SR networks can communicate across the LDP network:

- Configure Switch A, Switch B, Switch C, Switch D, and Switch E to run IS-IS.
- Configure Switch B, Switch C, and Switch D to run LDP.
- Configure Switch A, Switch B, Switch D, and Switch E to run MPLS SR.

**Figure 15 Network diagram**



**Table 46 Interface and IP address assignment**

| Device   | Interface   | IP address  | Device   | Interface   | IP address  |
|----------|-------------|-------------|----------|-------------|-------------|
| Switch A | Loop1       | 1.1.1.1/32  | Switch B | Loop1       | 2.2.2.2/32  |
|          | Vlan-Int100 | 10.0.0.1/24 |          | Vlan-Int100 | 10.0.0.2/24 |
| Switch C | Loop1       | 3.3.3.3/32  |          | Vlan-Int200 | 11.0.0.1/24 |
|          | Vlan-Int100 | 12.0.0.1/24 | Switch D | Loop1       | 4.4.4.4/32  |
|          | Vlan-Int200 | 11.0.0.2/24 |          | Vlan-Int100 | 12.0.0.2/24 |
| Switch E | Loop1       | 5.5.5.5/32  |          | Vlan-Int200 | 13.0.0.1/24 |
|          | Vlan-Int200 | 13.0.0.2/24 |          |             |             |

### Procedure

1. Configure IP addresses and masks for interfaces. (Details not shown.)
2. Configure IS-IS on the switches to achieve network level connectivity and set the IS-IS cost style to wide:

# Configure Switch A.

```
<SwitchA> system-view
[SwitchA] isis 1
[SwitchA-isis-1] network-entity 00.0000.0000.0001.00
[SwitchA-isis-1] cost-style wide
[SwitchA-isis-1] quit
[SwitchA] interface vlan-interface 100
[SwitchA-Vlan-interface100] isis enable 1
```

```

[SwitchA-Vlan-interface100] quit
[SwitchA] interface loopback 1
[SwitchA-LoopBack1] isis enable 1
[SwitchA-LoopBack1] quit
Configure Switch B.
<SwitchB> system-view
[SwitchB] isis 1
[SwitchB-isis-1] network-entity 00.0000.0000.0002.00
[SwitchB-isis-1] cost-style wide
[SwitchB-isis-1] quit
[SwitchB] interface vlan-interface 100
[SwitchB-Vlan-interface100] isis enable 1
[SwitchB-Vlan-interface100] quit
[SwitchB] interface vlan-interface 200
[SwitchB-Vlan-interface200] isis enable 1
[SwitchB-Vlan-interface200] quit
[SwitchB] interface loopback 1
[SwitchB-LoopBack1] isis enable 1
[SwitchB-LoopBack1] quit
Configure Switch C.
<SwitchC> system-view
[SwitchC] isis 1
[SwitchC-isis-1] network-entity 00.0000.0000.0003.00
[SwitchC-isis-1] cost-style wide
[SwitchC-isis-1] quit
[SwitchC] interface vlan-interface 100
[SwitchC-Vlan-interface100] isis enable 1
[SwitchC-Vlan-interface100] quit
[SwitchC] interface vlan-interface 200
[SwitchC-Vlan-interface200] isis enable 1
[SwitchC-Vlan-interface200] quit
[SwitchC] interface loopback 1
[SwitchC-LoopBack1] isis enable 1
[SwitchC-LoopBack1] quit
Configure Switch D.
<SwitchD> system-view
[SwitchD] isis 1
[SwitchD-isis-1] network-entity 00.0000.0000.0004.00
[SwitchD-isis-1] cost-style wide
[SwitchD-isis-1] quit
[SwitchD] interface vlan-interface 100
[SwitchD-Vlan-interface100] isis enable 1
[SwitchD-Vlan-interface100] quit
[SwitchD] interface vlan-interface 200
[SwitchD-Vlan-interface200] isis enable 1
[SwitchD-Vlan-interface200] quit
[SwitchD] interface loopback 1
[SwitchD-LoopBack1] isis enable 1

```

```
[SwitchD-LoopBack1] quit
Configure Switch E.
<SwitchE> system-view
[SwitchE] isis 1
[SwitchE-isis-1] network-entity 00.0000.0000.0005.00
[SwitchE-isis-1] cost-style wide
[SwitchE-isis-1] quit
[SwitchE] interface vlan-interface 200
[SwitchE-Vlan-interface200] isis enable 1
[SwitchE-Vlan-interface200] quit
[SwitchE] interface loopback 1
[SwitchE-LoopBack1] isis enable 1
[SwitchE-LoopBack1] quit
```

### 3. Configure LSR IDs on the switches:

```
Configure Switch A.
[SwitchA] mpls lsr-id 1.1.1.1
Configure Switch B.
[SwitchB] mpls lsr-id 2.2.2.2
Configure Switch C.
[SwitchC] mpls lsr-id 3.3.3.3
Configure Switch D.
[SwitchD] mpls lsr-id 4.4.4.4
Configure Switch E.
[SwitchE] mpls lsr-id 5.5.5.5
```

### 4. Configure LDP on Switch B, Switch C, and Switch D:

```
Configure Switch B.
[SwitchB] mpls ldp
[SwitchB-ldp] quit
[SwitchB] interface vlan-interface 200
[SwitchB-Vlan-interface200] mpls enable
[SwitchB-Vlan-interface200] mpls ldp enable
[SwitchB-Vlan-interface200] quit
Configure Switch C.
[SwitchC] mpls ldp
[SwitchC-ldp] quit
[SwitchC] interface vlan-interface 100
[SwitchC-Vlan-interface100] mpls enable
[SwitchC-Vlan-interface100] mpls ldp enable
[SwitchC-Vlan-interface100] quit
[SwitchC] interface vlan-interface 200
[SwitchC-Vlan-interface200] mpls enable
[SwitchC-Vlan-interface200] mpls ldp enable
[SwitchC-Vlan-interface200] quit
Configure Switch D.
[SwitchD] mpls ldp
[SwitchD-ldp] quit
[SwitchD] interface vlan-interface 100
```

```
[SwitchD-Vlan-interface1001] mpls enable
[SwitchD-Vlan-interface100] mpls ldp enable
[SwitchD-Vlan-interface100] quit
```

**5. Enable MPLS SR on Switch A, Switch B, Switch D, and Switch E:**

**# Configure Switch A.**

```
[SwitchA] isis 1
[SwitchA-isis-1] address-family ipv4
[SwitchA-isis-1-ipv4] segment-routing mpls
[SwitchA-isis-1-ipv4] quit
```

**# Configure Switch B.**

```
[SwitchB] isis 1
[SwitchB-isis-1] address-family ipv4
[SwitchB-isis-1-ipv4] segment-routing mpls
[SwitchB-isis-1-ipv4] quit
```

**# Configure Switch D.**

```
[SwitchD] isis 1
[SwitchD-isis-1] address-family ipv4
[SwitchD-isis-1-ipv4] segment-routing mpls
[SwitchD-isis-1-ipv4] quit
```

**# Configure Switch E.**

```
[SwitchE] isis 1
[SwitchE-isis-1] address-family ipv4
[SwitchE-isis-1-ipv4] segment-routing mpls
[SwitchE-isis-1-ipv4] quit
```

**6. Configure SRGBs on Switch A, Switch B, Switch D, and Switch E:**

**# Configure Switch A.**

```
[SwitchA-isis-1] segment-routing global-block 16000 16999
[SwitchA-isis-1] quit
```

**# Configure Switch B.**

```
[SwitchB-isis-1] segment-routing global-block 17000 17999
[SwitchB-isis-1] quit
```

**# Configure Switch D.**

```
[SwitchD-isis-1] segment-routing global-block 18000 18999
[SwitchD-isis-1] quit
```

**# Configure Switch E.**

```
[SwitchE-isis-1] segment-routing global-block 19000 19999
[SwitchE-isis-1] quit
```

**7. Configure IS-IS prefix SIDs for Switch A, Switch B, Switch D, and Switch E:**

**# Configure Switch A.**

```
[SwitchA] interface loopback 1
[SwitchA-LoopBack1] isis prefix-sid index 10
[SwitchA-LoopBack1] quit
```

**# Configure Switch B.**

```
[SwitchB] interface loopback 1
[SwitchB-LoopBack1] isis prefix-sid index 20
[SwitchB-LoopBack1] quit
```

**# Configure Switch D.**

```
[SwitchD] interface loopback 1
[SwitchD-LoopBack1] isis prefix-sid index 40
[SwitchD-LoopBack1] quit
Configure Switch E.
[SwitchE] interface loopback 1
[SwitchE-LoopBack1] isis prefix-sid index 50
[SwitchE-LoopBack1] quit
```

## Verifying the configuration

### # Display LDP LSP information on Switch B.

```
[SwitchB] display mpls ldp lsp
Status Flags: * - stale, L - liberal, B - backup, N/A - unavailable
FECs: 5 Ingress: 3 Transit: 3 Egress: 2
```

| FEC        | In/Out Label        | Nextthop | OutInterface |
|------------|---------------------|----------|--------------|
| 1.1.1.1/32 | 2171/-<br>-/2169(L) |          |              |
| 2.2.2.2/32 | 2175/-<br>-/2170(L) |          |              |
| 3.3.3.3/32 | -/2174              | 11.0.0.2 | Vlan200      |
|            | 2172/2174           | 11.0.0.2 | Vlan200      |
| 4.4.4.4/32 | -/2144              | 11.0.0.2 | Vlan200      |
|            | 2167/2144           | 11.0.0.2 | Vlan200      |
| 5.5.5.5/32 | -/2162              | 11.0.0.2 | Vlan200      |
|            | 2161/2162           | 11.0.0.2 | Vlan200      |

### # Display IS-IS SRLSP information on Switch B.

```
[SwitchB] display mpls lsp protocol isis
```

| FEC        | Proto | In/Out Label | Interface/Out NHLFE |
|------------|-------|--------------|---------------------|
| 1.1.1.1/32 | ISIS  | 17010/3      | Vlan100             |
| 1.1.1.1/32 | ISIS  | -/3          | Vlan100             |
| 2.2.2.2/32 | ISIS  | 17020/-      | -                   |
| 4.4.4.4/32 | ISIS  | 17040/2144   | Vlan200             |
| 4.4.4.4/32 | ISIS  | -/2144       | Vlan200             |
| 5.5.5.5/32 | ISIS  | 17050/2162   | Vlan200             |
| 5.5.5.5/32 | ISIS  | -/2162       | Vlan200             |

The output shows that the IS-IS SRLSP entries for Switch D and Switch E are using LDP labels.

## Command reference

### display isis segment-routing adjacency

Use **display isis segment-routing adjacency** to display IS-IS SR adjacency segment information.

#### Syntax

```
display isis segment-routing adjacency [process-id]
```

#### Views

Any view

## Predefined user roles

network-admin  
network-operator

## Parameters

*process-id*: Specifies an IS-IS process ID in the range of 1 to 65535. If you do not specify this argument, the command displays SR adjacency segment information about all IS-IS processes.

## Examples

# Display SR adjacency segment information about IS-IS process 1. (Applicable to routers.)

```
<Sysname> display isis segment-routing adjacency 1
```

```
Segment routing adjacency path information for IS-IS(1)
```

```

SystemID Interface Nexthop Label Type

0000.0000.0011 XGE1/0/1 10.123.110.21 24122 L1
0000.0000.0012 XGE1/0/2 192.168.56.21 24121 L1
```

# Display SR adjacency segment information about IS-IS process 1. (Applicable to switches.)

```
<Sysname> display isis segment-routing adjacency 1
```

```
Segment routing adjacency path information for IS-IS(1)
```

```

SystemID Interface Nexthop Label Type

0000.0000.0011 Vlan1 10.123.110.21 24122 L1
0000.0000.0012 Vlan2 192.168.56.21 24121 L1
```

**Table 47 Command output**

| Field     | Description                           |
|-----------|---------------------------------------|
| SystemID  | System ID of the neighbor.            |
| Interface | Interface connected to the neighbor.  |
| Nexthop   | Next hop of the adjacency segment.    |
| Label     | Incoming label.                       |
| Type      | Link adjacency type of the interface. |

## display isis segment-routing global-block

Use **display isis segment-routing global-block** to display IS-IS SRGB information.

## Syntax

```
display isis segment-routing global-block [level-1 | level-2]
[process-id]
```

## Views

Any view

## Predefined user roles

network-admin  
network-operator

## Parameters

**level-1**: Specifies Level-1 SRGB.

**level-2**: Specifies Level-2 SRGB.

*process-id*: Specifies an IS-IS process ID in the range of 1 to 65535. If you do not specify this argument, the command displays the SRGBs of all IS-IS processes.

## Usage guidelines

If you do not specify the **level-1** or **level-2** keyword, the command displays information about both the Level-1 SRGB and Level-2 SRGB.

## Examples

# Display information about the SRGBs for IS-IS process 1.

```
<Sysname> display isis segment-routing global-block 1
```

```
Segment routing global block information for IS-IS(1)

Level-1 SRGB

System ID Base Range

0000.0000.0011 16666 5557
0000.0000.0012 18012 4001
```

**Table 48 Command output**

| Field     | Description                      |
|-----------|----------------------------------|
| System ID | System ID of the neighbor.       |
| Base      | Minimum label value of the SRGB. |
| Range     | Number of labels in the SRGB.    |

## Related commands

**segment-routing global-block**

## isis prefix-sid

Use **isis prefix-sid** to configure an IS-IS prefix SID.

Use **undo isis prefix-sid** to restore the default.

## Syntax

```
isis prefix-sid { absolute absolute-value | index index-value }
[n-flag-clear] [explicit-null]
```

```
undo isis prefix-sid
```

## Default

No IS-IS prefix SID is configured.

## Views

Loopback interface view

## Predefined user roles

network-admin

## Parameters

**absolute** *absolute-value*: Specifies an absolute value as the prefix SID. The *absolute-value* argument is in the range of 0 to 1048575.

**index** *index-value*: Specifies an index value. The *index-value* argument is in the range of 0 to 1048575.

**n-flag-clear**: Sets the Node-SID flag bit of the prefix SID to 0 to use the prefix SID for a group of SR nodes. If you do not specify this keyword, the flag bit is 1 and the prefix SID is used for a single SR node.

**explicit-null**: Sets the Explicit-null flag bit of the prefix SID to 1. This setting requires that the upstream neighbor uses an explicit null label to replace the prefix SID. If you do not specify this keyword, the flag bit is 0 and the upstream neighbor continues to forward the packet based on the prefix SID. For more information about the explicit null label, see MPLS basics configuration in *MPLS Configuration Guide*.

## Usage guidelines

An absolute value used as the prefix SID must be in the SRGB of the node.

If you specify an index value, the sum of the index value and the SRGB base value is used as the prefix SID. The prefix SID must be in the SRGB of the node.

To use a prefix SID for a group of SR nodes in anycast scenarios, specify the **n-flag-clear** keyword to set the Node-SID flag bit of the prefix SID to 0.

To configure an IS-IS prefix SID, you must enable an IS-IS process on the loopback interface.

## Examples

# Configure absolute value 20 as the IS-IS prefix SID on loopback interface 1.

```
<Sysname> system-view
[Sysname] interface loopback 1
[Sysname-LoopBack1] isis enable 1
[Sysname-LoopBack1] isis prefix-sid index 20
```

## segment-routing adjacency enable

Use **segment-routing adjacency enable** to enable MPLS SR adjacency label allocation.

Use **undo segment-routing adjacency enable** to disable MPLS SR adjacency label allocation.

## Syntax

```
segment-routing adjacency enable
undo segment-routing adjacency enable
```

## Default

MPLS SR adjacency label allocation is disabled.

## Views

IS-IS IPv4 unicast address family view

## Predefined user roles

network-admin

## Usage guidelines

For this command to take effect, you must enable MPLS SR.

## Examples

```
Enable MPLS SR adjacency label allocation for IS-IS process 1.
<Sysname> system-view
[Sysname] isis 1
[Sysname-isis-1] address-family ipv4
[Sysname-isis-1-ipv4] segment-routing adjacency enable
```

## Related commands

**segment-routing mpls**

## segment-routing global-block

Use **segment-routing global-block** to configure the MPLS SRGB.

Use **undo segment-routing global-block** to restore the default.

## Syntax

```
segment-routing global-block minimum-value maximum-value
undo segment-routing global-block
```

## Default

The minimum label value is 16000. The maximum label value is 24000.

## Views

IS-IS view

## Predefined user roles

network-admin

## Parameters

*minimum-value*: Specifies the minimum label value, in the range of 16000 to 23999.

*maximum-value*: Specifies the maximum label value, in the range of 16001 to 24000.

## Usage guidelines

To configure the SRGB successfully, make sure the SRGB contains the configured prefix SID.

## Examples

```
Configure the MPLS SRGB to be from 17000 to 22000 for IS-IS process 1.
<Sysname> system-view
[Sysname] isis 1
[Sysname-isis-1] segment-routing global-block 17000 22000
```

## segment-routing mpls

Use **segment-routing mpls** to enable MPLS SR.

Use **undo segment-routing mpls** to disable MPLS SR.

## Syntax

```
segment-routing mpls
undo segment-routing mpls
```

## Default

MPLS SR is disabled.

## Views

IS-IS IPv4 unicast address family view

## Predefined user roles

network-admin

## Usage guidelines

For MPLS SR to take effect, set the IS-IS cost style to wide, compatible, or wide-compatible before configuring IS-IS to support MPLS SR. For more information about the cost style, see IS-IS configuration in *Layer 3—IP Routing Configuration Guide*.

## Examples

```
Enable MPLS SR in IS-IS IPv4 unicast address family view of IS-IS process 1.
<Sysname> system-view
[Sysname] isis 1
[Sysname-isis-1] cost-style wide
[Sysname-isis-1] address-family ipv4
[Sysname-isis-1-ipv4] segment-routing mpls
```

## Related commands

**cost-style** (*Layer 3—IP Routing Command Reference*)

## segment-routing sr-prefer

Use **segment-routing sr-prefer** to configure the device to prefer SRLSPs in traffic forwarding.

Use **undo segment-routing sr-prefer** to restore the default.

## Syntax

```
segment-routing sr-prefer [prefix-list prefix-list-name]
undo segment-routing sr-prefer
```

## Default

The device prefers LDP LSPs in traffic forwarding.

## Views

IS-IS IPv4 unicast address family view

## Predefined user roles

network-admin

## Parameters

**prefix-list prefix-list-name**: Specifies an IP address prefix list by its name, a case-sensitive string of 1 to 63 characters. The device preferentially uses SRLSPs to forward traffic destined for addresses on the prefix list, and preferentially uses LDP LSPs to forward traffic destined for other addresses. If you do not specify an IP address prefix list, the device preferentially uses SRLSPs to forward all traffic.

## Usage guidelines

This command determines whether the device prefers SRLSPs or LDP LSPs when both SRLSPs and LDP LSPs are available for traffic forwarding.

This command takes effect only when MPLS SR is enabled and the SRLSPs use prefix SIDs.

## Examples

# Configure the device to preferentially use SRLSPs established by OSPF to forward all traffic.

```
<Sysname> system-view
[Sysname] ospf 1
[Sysname-ospf-1] segment-routing sr-prefer
```

# Configure the device to preferentially use SRLSPs established by OSPF to forward traffic that is identified by prefix list 8.

```
<Sysname> system-view
[Sysname] ip prefix-list 8 permit 4.4.4.4 32
[Sysname] ospf 1
[Sysname-ospf-1] segment-routing sr-prefer prefix-list 8
```

## Related commands

**ip prefix-list** (*Layer 3—IP Routing Command Reference*)

**segment-routing mpls**

# New feature: Displaying buffer usage statistics for interfaces

## Displaying buffer usage statistics for interfaces

Execute **display** commands in any view.

| Task                                            | Command                                                                                                                                                          |
|-------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Display buffer usage statistics for interfaces. | <b>display buffer usage interface</b> [ <i>interface-type</i> [ <i>interface-number</i> ] ] [ <b>queue</b> <i>queue-id</i> <b>verbose</b> <i>number-record</i> ] |

## Command reference

### display buffer usage interface

Use **display buffer usage interface** to display buffer usage statistics for interfaces.

#### Syntax

```
display buffer usage interface [interface-type [interface-number]] [queue queue-id verbose number-record]
```

#### Views

Any view

#### Predefined user roles

network-admin

network-operator

## Parameters

*interface-type* [*interface-number*]: Specifies an interface by its type and number. If you do not specify the *interface-type* argument, this command displays buffer usage statistics for all Ethernet interfaces. If you specify the *interface-type* argument without the *interface-number* argument, this command displays buffer usage statistics for all Ethernet interfaces of the specified type.

**queue** *queue-id* **verbose** *number-record*: Displays buffer usage threshold-exceeding statistics for a queue, including threshold-exceeding start time, threshold-exceeding duration, and number of bytes forwarded by the queue within the threshold-exceeding duration. The value range for the *queue-id* argument is 0 to 7. The *number-record* argument represents the maximum number of threshold violations to be displayed, in the range of 1 to 2000. If you do not specify this option, the command displays buffer usage statistics. If a queue does not exceed the buffer usage threshold and you specify this option, the command displays nothing.

## Examples

# Display buffer usage statistics for Ten-GigabitEthernet 1/0/1.

```
<Sysname> display buffer usage interface ten-gigabitethernet 1/0/1
```

| Interface | QueueID | Total   | Used | Threshold(%) | Violations |
|-----------|---------|---------|------|--------------|------------|
| -----     |         |         |      |              |            |
| XGE1/0/1  | 0       | 9418032 | 0    | 30           | 0          |
|           | 1       | 9418032 | 0    | 30           | 0          |
|           | 2       | 9418032 | 0    | 30           | 0          |
|           | 3       | 9418032 | 0    | 30           | 0          |
|           | 4       | 9418032 | 0    | 30           | 0          |
|           | 5       | 9418032 | 0    | 30           | 0          |
|           | 6       | 9418032 | 0    | 30           | 0          |
|           | 7       | 9418032 | 0    | 30           | 0          |

# Display buffer usage threshold-exceeding statistics for queue 1 on Ten-GigabitEthernet 1/0/1.

```
<Sysname> display buffer usage interface ten-gigabitethernet 1/0/1 queue 1 verbose 10
```

| Start Time              | Continuous(ms) | Used   |
|-------------------------|----------------|--------|
| -----                   |                |        |
| 2011/03/14 07:51:57:525 | 20531          | 822224 |
| 2011/03/14 07:49:25:361 | 70535          | 826384 |

**Table 49 Command output**

| Field          | Description                                                                                               |
|----------------|-----------------------------------------------------------------------------------------------------------|
| Total          | Data buffer size in bytes allowed for a queue.                                                            |
| Used           | Data buffer size in bytes that has been used by a queue.                                                  |
| Threshold(%)   | Buffer usage threshold for a queue. The threshold value is the same as the per-interface threshold value. |
| Violations     | Number of threshold violations for a queue.<br>The value of this field is reset upon a switch reboot.     |
| Start Time     | Threshold-exceeding start time.                                                                           |
| Continuous(ms) | Threshold-exceeding duration in milliseconds.                                                             |

# New feature: Configuring an EAP profile

## Configuring an EAP profile

### About EAP profiles

An EAP profile is a collection of EAP authentication settings, including the EAP authentication method and the CA certificate file to be used for some EAP authentication methods.

### Restrictions and guidelines

You can specify an EAP profile in multiple test profiles.

You can configure a maximum of 16 EAP profiles.

### Prerequisites

Before you specify a CA certificate file, use FTP or TFTP to transfer the CA certificate file to the root directory of the default storage medium on the device.

In an IRF fabric, make sure a CA certificate file already exists in the root directory of the default storage medium on the master device before you specify the file.

### Procedure

1. Enter system view.  
**system-view**
2. Create an EAP profile and enter EAP profile view.  
**eap-profile** *eap-profile-name*
3. Specify the EAP authentication method.  
**method** { **md5** | **peap-gtc** | **peap-mschapv2** | **ttls-gtc** | **ttls-mschapv2** }  
By default, the EAP authentication method is MD5-challenge.
4. Specify a CA certificate file for EAP authentication.  
**ca-file** *file-name*  
By default, no CA certificate file is specified for EAP authentication.  
You must specify a CA certificate file to verify the RADIUS server certificate if the EAP authentication method is PEAP-GTC, PEAP-MSCHAPv2, TTLS-GTC, or TTLS-MSCHAPv2.

## Command reference

### eap-profile

Use **eap-profile** to create an EAP profile and enter its view, or enter the view of an existing EAP profile.

Use **undo eap-profile** to delete an EAP profile.

### Syntax

**eap-profile** *eap-profile-name*  
**undo eap-profile** *eap-profile-name*

### Default

No EAP profiles exist.

## Views

System view

## Predefined user roles

network-admin

## Parameters

*eap-profile-name*: Specifies the EAP profile name, a case-sensitive string of 1 to 32 characters.

## Usage guidelines

An EAP profile is a collection of EAP authentication settings, including the EAP authentication method and the CA certificate file to be used for some EAP authentication methods. You can use an EAP profile in a test profile for RADIUS server status detection.

You can specify an EAP profile in multiple test profiles.

You can configure a maximum of 16 EAP profiles.

## Examples

# Create an EAP profile named **eap1** and enter its view.

```
<Sysname> system-view
[Sysname] eap-profile eap1
[Sysname-eap-profile-eap1]
```

## Related commands

**radius-server test-profile**

## ca-file

Use **ca-file** to specify a CA certificate file for EAP authentication.

Use **undo ca-file** to restore the default.

## Syntax

```
ca-file file-name
undo ca-file
```

## Default

No CA certificate file is specified for EAP authentication. The device does not verify the RADIUS server certificate during EAP authentication.

## Views

EAP profile view

## Predefined user roles

network-admin

## Parameters

*file-name*: Specifies a CA certificate file by its name, a case-sensitive string of 1 to 91 characters.

## Usage guidelines

You must specify a CA certificate file to verify the RADIUS server certificate if the EAP authentication method is PEAP-GTC, PEAP-MSCHAPv2, TTLS-GTC, or TTLS-MSCHAPv2.

Before you specify a CA certificate file, you must use FTP or TFTP to transfer the CA certificate file to the root directory of the default storage medium on the device.

In an IRF fabric, make sure a CA certificate file already exists in the root directory of the default storage medium on the master device before you specify the file.

You can specify only one CA certificate file in an EAP profile. If you execute this command multiple times, the most recent configuration takes effect.

If you change the CA certificate file, the new CA certificate file takes effect at the next server status detection.

## Examples

# In EAP profile **eap1**, specify CA certificate file **CA.der** for EAP authentication.

```
<Sysname> system-view
[Sysname] eap-profile eap1
[Sysname-eap-profile-eap1] ca-file CA.der
```

## method

Use **method** to specify the EAP authentication method.

Use **undo method** to restore the default.

## Syntax

```
method { md5 | peap-gtc | peap-mschapv2 | ttls-gtc | ttls-mschapv2 }
undo method
```

## Default

MD5-challenge authentication is used.

## Views

EAP profile view

## Predefined user roles

network-admin

## Parameters

**md5**: Specifies the MD5-challenge method.

**peap-gtc**: Specifies the PEAP-GTC method.

**peap-mschapv2**: Specifies the PEAP-MSCHAPv2 method.

**ttls-gtc**: Specifies the TTLS-GTC method.

**ttls-mschapv2**: Specifies the TTLS-MSCHAPv2 method.

## Usage guidelines

You must specify an EAP authentication method that is supported by the RADIUS server to be detected.

You can specify only one EAP authentication method in an EAP profile. If you execute this command multiple times, the most recent configuration takes effect.

If you change the EAP authentication method, the new method takes effect in the next server status detection.

## Examples

# In EAP profile **eap1**, specify PEAP-GTC as the EAP authentication method.

```
<Sysname> system-view
[Sysname] eap-profile eap1
[Sysname-eap-profile-eap1] method peap-gtc
```

# New feature: AAA test

## Configuring the AAA test feature

### About the AAA test feature

This feature enables the device to send authentication or accounting requests to the specified AAA servers to simulate an authentication or accounting process of a user. Use this feature to identify the reasons for the failure of the interaction between the device and the AAA servers. This feature is applicable only to RADIUS.

When performing an AAA test, the device ignores the status of the specified AAA servers and the RADIUS server load sharing feature. The process of an AAA test is as follows:

1. The device sends authentication requests that carry the specified username and password to the specified authentication server or to the authentication servers in the specified RADIUS scheme. The device tries to communicate with the authentication servers in the specified scheme in sequence.

The process goes to the next step in the following situations:

- The device receives an authentication response (no matter the authentication succeeds or fails).
- The device does not receive any authentication response after making all authentication request attempts.

This step is skipped if no correct authentication server is specified for the AAA test or no authentication servers are configured in the specified RADIUS scheme.

2. The device sends start-accounting requests to the specified accounting server or to the accounting servers in the specified RADIUS scheme. The device tries to communicate with the accounting servers in the specified scheme in sequence.

The process goes to the next step in the following situations:

- The device receives a start-accounting response (no matter the accounting succeeds or fails).
- The device does not receive any start-accounting response after making all start-accounting request attempts.

This step and the next step are skipped if no correct accounting server is specified for the AAA test or no accounting servers are configured in the specified RADIUS scheme.

3. The device sends stop-accounting requests to the accounting servers to which it has sent a start-accounting request.

The process finishes in the following situations:

- The device receives a stop-accounting response.
- The device does not receive any stop-accounting response after making all stop-accounting request attempts.

To identify attributes that cause authentication or accounting failures, you can configure the device to carry specific attributes in RADIUS requests or define values for specific attributes in the requests. [Table 50](#) shows the attributes that RADIUS requests carry by default.

**Table 50 Attributes that RADIUS requests carry by default**

| Packet type                   | Attributes that the type of packets carry by default                                                    |
|-------------------------------|---------------------------------------------------------------------------------------------------------|
| RADIUS authentication request | User-Name<br>CHAP-Password (or User-Password)<br>CHAP-Challenge<br>NAS-IP-Address (or NAS-IPv6-Address) |

| Packet type               | Attributes that the type of packets carry by default                                                                                                  |
|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
|                           | Service-Type<br>Framed-Protocol<br>NAS-Identifier<br>NAS-Port-Type<br>Acct-Session-Id                                                                 |
| RADIUS accounting request | User-Name<br>Acct-Status-Type<br>NAS-IP-Address (or NAS-IPv6-Address)<br>NAS-Identifier<br>Acct-Session-Id<br>Acct-Delay-Time<br>Acct-Terminate-Cause |

## Restrictions and guidelines

When you perform an AAA test, follow these restrictions and guidelines:

- The device might communicate with the AAA servers incorrectly during an AAA test. Make sure no users come online or go offline during an AAA test.
- If the configuration of the specified RADIUS scheme changes, the new configuration does not affect the current AAA test. The modification will take effect in the next test.
- The system can have only one AAA test at a time. Another AAA test can be performed only after the current test finishes.

When you configure attributes to be included in or excluded from RADIUS requests, follow these restrictions and guidelines:

- Before you include an attribute that is already configured to be excluded from RADIUS requests, you must cancel the exclusion configuration by using the **undo exclude** command.
- Before you exclude an attribute that is already configured to be included in RADIUS requests, you must cancel the inclusion configuration by using the **undo include** command.

## Prerequisites

Before you perform an AAA test, you must configure a RADIUS scheme that contains the RADIUS servers to be tested.

Plan the RADIUS attributes to be included in RADIUS requests. Besides the attributes carried by default, the device adds the specified attributes to RADIUS packets in the order that they are specified by using the **include** command. Additional attributes cannot be added to a RADIUS request if the length of the RADIUS request reaches 4096 bytes.

## Procedure

1. (Optional.) Configure a RADIUS attribute test group:
  - a. Enter system view.  
**system-view**
  - b. Create a RADIUS attribute test group and enter its view.  
**radius attribute-test-group** *attr-test-group-name*  
You can create multiple RADIUS attribute test groups.
  - c. Include an attribute in RADIUS requests.  
**include** { **accounting** | **authentication** } { **name** *attribute-name* | [ **vendor** *vendor-id* ] **code** *attribute-code* } **type** { **binary** | **date** |

```
integer | interface-id | ip | ipv6 | ipv6-prefix | octets | string }
value attribute-value
```

Use this command to add attributes that RADIUS requests do not carry by default to the RADIUS requests.

For an attribute that RADIUS requests carry by default, you can use this command to change its attribute value.

- d. Exclude an attribute from RADIUS requests.

```
exclude { accounting | authentication } name attribute-name
```

Use this command to exclude an attribute that RADIUS requests carry by default from the RADIUS requests sent during an AAA test to help troubleshoot authentication or accounting failures.

- e. Return to system view.

```
quit
```

- f. Return to user view.

```
quit
```

- 2. Perform an AAA test in user view.

```
test-aaa user user-name password password radius-scheme
radius-scheme-name [radius-server { ipv4-address | ipv6
ipv6-address } port-number [vpn-instance vpn-instance-name]] [chap
| pap] [attribute-test-group attr-test-group-name] [trace]
```

## Command reference

### exclude

Use **exclude** to exclude an attribute from RADIUS requests.

Use **undo exclude** to cancel the configuration of excluding an attribute from RADIUS requests.

#### Syntax

```
exclude { accounting | authentication } name attribute-name
undo exclude { accounting | authentication } name attribute-name
```

#### Default

No attributes are configured to be excluded from RADIUS requests.

#### Views

RADIUS attribute test group view

#### Predefined user roles

network-admin

#### Parameters

**accounting**: Specifies RADIUS accounting requests.

**authentication**: Specifies RADIUS authentication requests.

**name** *attribute-name*: Specifies a RADIUS attribute by its name, a case-insensitive string of 1 to 63 characters. The specified attribute must be an attribute that RADIUS requests carry by default. RADIUS authentication requests carry the following attributes by default: Service-Type, Framed-Protocol, NAS-Identifier, Acct-Session-Id, and NAS-Port-Type. RADIUS accounting requests carry the following attributes by default: NAS-Identifier, Acct-Delay-Time, Acct-Session-Id, and Acct-Terminate-Cause.

## Usage guidelines

Use this command to exclude an attribute from RADIUS requests sent during an AAA test to help troubleshoot authentication or accounting failures.

Before you exclude an attribute that is already configured to be included in RADIUS requests, you must cancel the inclusion configuration by using the **undo include** command.

## Examples

# In RADIUS attribute test group **t1**, exclude Service-Type attribute from RADIUS authentication requests.

```
<Sysname> system-view
```

```
[Sysname] radius attribute-test-group t1
```

```
[Sysname-radius-attr-test-grp-t1] exclude authentication name Service-Type
```

## Related commands

**include**

**test-aaa**

## include

Use **include** to include an attribute in RADIUS requests.

Use **undo include** to cancel the configuration of including an attribute in RADIUS requests.

## Syntax

```
include { accounting | authentication } { name attribute-name | [vendor vendor-id] code attribute-code } type { binary | date | integer | interface-id | ip | ipv6 | ipv6-prefix | octets | string } value attribute-value
```

```
undo include { accounting | authentication } { name attribute-name | [vendor vendor-id] code attribute-code }
```

## Default

No attributes are configured to be included in RADIUS authentication or accounting requests.

## Views

RADIUS attribute test group view

## Predefined user roles

network-admin

## Parameters

**accounting**: Specifies RADIUS accounting requests.

**authentication**: Specifies RADIUS authentication requests.

**name** *attribute-name*: Specifies a standard RADIUS attribute by its name, a case-insensitive string of 1 to 63 characters.

**vendor** *vendor-id*: Specifies a vendor by its ID in the range of 1 to 65535. If you do not specify a vendor, this command includes a standard attribute in RADIUS requests.

**code** *attribute-code*: Specifies a RADIUS attribute by its code in the range of 1 to 255.

**type**: Specifies a data type for the attribute content.

**binary**: Binary type.

**date**: Date type.

**integer**: Integer type.

**interface-id**: Interface ID type.

**ip**: IPv4 address type.

**ipv6**: IPv6 address type.

**ipv6-prefix**: IPv6 address prefix type.

**octets**: Octet type.

**string**: String type.

**value attribute-value**: Specifies the value for the attribute of the data type. The value range of the *attribute-value* argument varies by data type.

- For the binary type, the value is a string of 1 to 256 hexadecimal characters, which represents a binary number with a maximum of 128 bytes.
- For the date type, the value range is 0 to 4294967295.
- For the integer type, the value range is 0 to 4294967295.
- For the interface ID type, the value range is 1 to ffffffffffff.
- For the IPv6 address prefix type, the value is in the format of *prefix/prefix-length*.
- For the octet type, the value is a string of 1 to 256 hexadecimal characters, which represents an octet number with a maximum of 128 bytes.
- For the string type, the value of this argument is a string of 1 to 253 characters.

## Usage guidelines

Use this command to add an attribute that RADIUS requests do not carry by default to the RADIUS requests. The **undo** form of this command removes the attribute from the RADIUS requests.

For an attribute that RADIUS requests carry by default, you can use this command to change its value. The **undo** form of this command restores the attribute value to the default.

Table 51 shows the attributes that RADIUS requests carry by default.

**Table 51 Attributes that RADIUS requests carry by default**

| Packet type                   | Attributes that the type of packets carry by default                                                                                                                                  |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RADIUS authentication request | User-Name, CHAP-Password (or User-Password), CHAP-Challenge, NAS-IP-Address (or NAS-IPv6-Address), Service-Type, Framed-Protocol, NAS-Identifier, NAS-Port-Type, and Acct-Session-Id. |
| RADIUS accounting request     | User-Name, Acct-Status-Type, NAS-IP-Address (or NAS-IPv6-Address), NAS-Identifier, Acct-Session-Id, Acct-Delay-Time, and Acct-Terminate-Cause.                                        |

For the accuracy of AAA tests, the value of an attribute must be of the data type specified for that attribute.

The attribute names of standard attributes saved in the configuration file will be converted to attribute codes.

Before you include an attribute that is already configured to be excluded from RADIUS requests, you must cancel the exclusion configuration by using the **undo exclude** command.

Plan the RADIUS attributes to be included in RADIUS requests. Besides the attributes carried by default, the device adds the specified attributes to RADIUS packets in the order that they are specified by using the **include** command. Additional attributes cannot be added to a RADIUS request if the length of the RADIUS request reaches 4096 bytes.

## Examples

```
In RADIUS attribute test group t1, include Calling-Station-Id attribute with value
08-00-27-00-34-D8 in RADIUS authentication requests.
<Sysname> system-view
[Sysname] radius attribute-test-group t1
[Sysname-radius-attr-test-grp-t1] include authentication name Calling-Station-Id type
string value 08-00-27-00-34-d8
```

## Related commands

```
exclude
test-aaa
```

## radius attribute-test-group

Use **radius attribute-test-group** to create a RADIUS attribute test group and enter its view, or enter the view of an existing RADIUS attribute test group.

Use **undo radius attribute-test-group** to remove a RADIUS attribute test group.

## Syntax

```
radius attribute-test-group attr-test-group-name
undo radius attribute-test-group attr-test-group-name
```

## Default

No RADIUS attribute test groups exist.

## Views

System view

## Predefined user roles

network-admin

## Parameters

*attr-test-group-name*: Specifies the name of a RADIUS attribute test group, a case-insensitive string of 1 to 31 characters.

## Usage guidelines

A RADIUS attribute test group is a collection of RADIUS attributes that will be included in or excluded from RADIUS requests.

The system can have multiple RADIUS attribute test groups.

## Examples

```
Create a RADIUS attribute test group named t1 and enter its view.
<Sysname> system-view
[Sysname] radius attribute-test-group t1
[Sysname-radius-attr-test-grp-t1]
```

## Related commands

```
exclude
include
test-aaa
```

## test-aaa

Use **test-aaa** to perform an AAA test.

### Syntax

```
test-aaa user user-name password password radius-scheme
radius-scheme-name [radius-server { ipv4-address | ipv6 ipv6-address }
port-number [vpn-instance vpn-instance-name]] [chap | pap]
[attribute-test-group attr-test-group-name] [trace]
```

### Views

User view

### Predefined user roles

network-admin

### Parameters

**user** *user-name*: Specifies the test username, a string of 1 to 80 characters. The username can be a pure username or contain a domain name. The format for a username containing a domain name is *pure-username@domain-name*. The pure username is case sensitive and the domain name is case insensitive.

**password** *password*: Specifies the password of the test user, a case-sensitive string of 1 to 63 characters.

**radius-scheme** *radius-scheme-name*: Specifies a RADIUS scheme by its name, a case-insensitive string of 1 to 32 characters.

**radius-server**: Specifies a RADIUS server.

*ipv4-address*: Specifies the IPv4 address of the RADIUS server.

**ipv6** *ipv6-address*: Specifies the IPv6 address of the RADIUS server.

*port-number*: Specifies the UDP port number of the RADIUS server, in the range of 1 to 65535.

**vpn-instance** *vpn-instance-name*: Specifies the MPLS L3VPN instance to which the RADIUS server belongs. The *vpn-instance-name* argument is a case-sensitive string of 1 to 31 characters. If the server is on the public network, do not specify this option.

**chap**: Specifies the CHAP authentication method (the default).

**pap**: Specifies the PAP authentication method.

**attribute-test-group** *attr-test-group-name*: Specifies a RADIUS attribute test group by its name, a case-insensitive string of 1 to 31 characters. If you do not specify a RADIUS attribute test group or the specified RADIUS attribute test group does not exist, the device does not change the attributes carried in authentication or accounting requests.

**trace**: Displays detailed information about RADIUS packets exchanged during the AAA test. If you do not specify this keyword, the command displays brief information about the AAA test, including the sent and received packets and the test result.

### Usage guidelines

Use this command to identify the reasons for the failure of interaction between the device and the AAA servers.

The device might communicate with the AAA servers incorrectly during an AAA test. Make sure no users come online or go offline during an AAA test.

If the configuration of the specified RADIUS scheme changes, the new configuration does not affect the current AAA test. The modification will take effect in the next test.

The system can have only one AAA test at a time. Another AAA test can be performed only after the current test finishes.

## Examples

# Perform an AAA test and display detailed information about the test. The test uses username **user1**, password **123456**, the CHAP authentication method, and RADIUS scheme **test**.

```
<Sysname> test-aaa user user1 password 123456 radius-scheme test chap trace
```

Sent a RADIUS authentication request.

```
Server IP : 192.168.1.110
Source IP : 192.168.1.166
VPN instance : N/A
Server port : 1812
Packet type : Authentication request
Packet length : 118 bytes
Packet ID : 0
Attribute list:
[User-Name(1)] [6] [user1]
[CHAP-Password(3)] [19] [*****]
[NAS-IP-Address(4)] [6] [192.168.1.166]
[Service-Type(6)] [6] [2] [Framed]
[Framed-Protocol(7)] [6] [1] [PPP]
[NAS-Identifier(32)] [5] [Sysname]
[Acct-Session-Id(44)] [40] [00000008201707241008280000000c16100171]
[CHAP-Challenge(60)] [18] [*****]
[NAS-Port-Type(61)] [6] [15] [Ethernet]
```

Received a RADIUS authentication response.

```
Server IP : 192.168.1.110
Source IP : 192.168.1.166
VPN instance : N/A
Server port : 1812
Packet type : Access-Reject
Packet length : 20 bytes
Packet ID : 0
Reply-Message: "E63032: Incorrect password. You can retry 9 times."
```

Sent a RADIUS start-accounting request.

```
Server IP : 192.168.1.110
Source IP : 192.168.1.166
VPN instance : N/A
Server port : 1813
Packet type : Start-accounting request
Packet length : 63 bytes
Packet ID : 1
Attribute list:
[User-Name(1)] [6] [user1]
[Acct-Status-Type(40)] [6] [1] [Start]
[NAS-IP-Address(4)] [6] [192.168.1.166]
[NAS-Identifier(32)] [5] [Sysname]
```

```

[Acct-Session-Id(44)] [40] [00000008201707241008280000000c16100171]

Received a RADIUS start-accounting response.
Server IP : 192.168.1.110
Source IP : 192.168.1.166
VPN instance : N/A
Server port : 1813
Packet type : Start-accounting response
Packet length: 20 bytes
Packet ID : 1

Sent a RADIUS stop-accounting request.
Server IP : 192.168.1.110
Source IP : 192.168.1.166
VPN instance : N/A
Server port : 1813
Packet type : Stop-accounting request
Packet length: 91 bytes
Packet ID : 1
Attribute list:
[User-Name(1)] [6] [user1]
[Acct-Status-Type(40)] [6] [2] [Stop]
[NAS-IP-Address(4)] [6] [192.168.1.166]
[NAS-Identifier(32)] [5] [Sysname]
[Acct-Delay-Time(41)] [6] [0]
[Acct-Session-Id(44)] [40] [00000008201707241008280000000c16100171]
[Acct-Terminate-Cause(49)] [6] [1] [User Request]

Received a RADIUS stop-accounting response.
Server IP : 192.168.1.110
Source IP : 192.168.1.166
VPN instance : N/A
Server port : 1813
Packet type : Stop-accounting response
Packet length: 20 bytes
Packet ID : 1

Test result: Failed

Perform an AAA test and display brief information about the test. The test uses username user1,
password 123456 and the CHAP authentication method to test RADIUS server at 192.168.1.110 in
RADIUS scheme test.

<Sysname> test-aaa user user1 password 123456 radius-scheme test radius-server
192.168.1.110 1812
Sent a RADIUS authentication request.
Received a RADIUS authentication response.

Test result: Successful

```

**Table 52 Command output**

| Field                                                                                         | Description                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Server IP                                                                                     | IP address of the server.                                                                                                                                                                                                                                                                                                                                    |
| Source IP                                                                                     | Source IP address of the RADIUS packet.                                                                                                                                                                                                                                                                                                                      |
| VPN instance                                                                                  | MPLS L3VPN instance to which the server belongs. This field displays <b>N/A</b> if the server belongs to the public network.                                                                                                                                                                                                                                 |
| Server port                                                                                   | UDP port number of the server.                                                                                                                                                                                                                                                                                                                               |
| Packet type                                                                                   | Type of the RADIUS packet: <ul style="list-style-type: none"> <li>• Authentication request.</li> <li>• Access-Accept.</li> <li>• Access-Reject.</li> <li>• Start-accounting request.</li> <li>• Start-accounting response.</li> <li>• Stop-accounting request.</li> <li>• Stop-accounting response.</li> </ul>                                               |
| Packet length                                                                                 | Total length of the RADIUS packet, in bytes.                                                                                                                                                                                                                                                                                                                 |
| Packet ID                                                                                     | ID of the RADIUS packet. This field is used to identity a pair of request and response packets.                                                                                                                                                                                                                                                              |
| [ <i>attribute-name (code)</i> ] [ <i>length</i> ]<br>[ <i>value</i> ] [ <i>description</i> ] | Information about a RADIUS attribute: <ul style="list-style-type: none"> <li>• <b>attribute-name</b>—Name of the attribute.</li> <li>• <b>code</b>—Code of the attribute.</li> <li>• <b>length</b>—Length of the attribute, in bytes.</li> <li>• <b>value</b>—Value of the attribute.</li> <li>• <b>description</b>—Description of the attribute.</li> </ul> |
| Reply-Message:                                                                                | The RADIUS server rejected the authentication request and replied a message.                                                                                                                                                                                                                                                                                 |
| Test result                                                                                   | Result of the AAA test: <ul style="list-style-type: none"> <li>• <b>Successful</b>—The test has succeeded.</li> <li>• <b>Failed</b>—The test has failed. If any request is rejected, the test fails.</li> </ul>                                                                                                                                              |

## Related commands

`radius scheme`

`radius attribute-test-group`

## New feature: Displaying and clearing history authentication and accounting load statistics for all RADIUS servers

### Displaying and clearing history authentication and accounting load statistics for all RADIUS servers

Execute `display` commands in any view and `reset` commands in user view.

| Task                                                                                | Command                                            |
|-------------------------------------------------------------------------------------|----------------------------------------------------|
| Display authentication and accounting load statistics for all RADIUS servers.       | <code>display radius server-load statistics</code> |
| Clear history authentication and accounting load statistics for all RADIUS servers. | <code>reset radius server-load statistics</code>   |

## Command reference

### display radius server-load statistics

Use `display radius server-load statistics` to display authentication and accounting load statistics for all RADIUS servers.

#### Syntax

```
display radius server-load statistics
```

#### Views

Any view

#### Predefined user roles

network-admin  
network-operator

#### Usage guidelines

This command displays the following statistics:

- **Last-5-second statistics**—Total number of authentication or accounting requests sent to each RADIUS server in the last 5 seconds.
- **History statistics**—Total number of authentication or accounting requests sent to each RADIUS server since the device starts up.

The device collects the statistics as follows:

- **Last-5-second statistics**—From the device sends the first authentication or accounting request to a RADIUS server, the device counts the number of authentication or accounting requests sent to the server every 5 seconds. Then, the device updates the last-5-second authentication and accounting statistics for the server.
- **History statistics**—The device increases the history statistics for a RADIUS server by 1 each time it sends an authentication or accounting request to the server. The device does not decrease the history statistics even though users go offline or the server fails to response to a request within the timeout time.

Based on the statistics, you can adjust the load on RADIUS servers by changing the sequence in which the servers are configured or the weight values of the servers.

This command displays statistics only for RADIUS servers whose IP addresses are available or can be resolved from their hostnames.

The device deletes all statistics for a RADIUS server if the server is removed from a RADIUS scheme or the server's IP address, VPN instance, or service port number changes.

If an active/standby switchover occurs, the last-5-second statistics are deleted. However, the history statistics are not deleted. The history statistics might be inaccurate.

If the device reboots, both the last-5-seconds statistics and the history statistics are deleted.

## Examples

# Display authentication and accounting load statistics for all RADIUS servers.

```
<Sysname> display radius server-load statistics
```

Authentication servers: 2

| IP      | VPN | Port | Last 5 sec | History |
|---------|-----|------|------------|---------|
| 1.1.1.1 | N/A | 1812 | 20         | 100     |
| 2.2.2.2 | ABC | 1812 | 0          | 20      |

Accounting servers: 2

| IP      | VPN | Port | Last 5 sec | History |
|---------|-----|------|------------|---------|
| 1.1.1.1 | N/A | 1813 | 20         | 100     |
| 2.2.2.2 | ABC | 1813 | 0          | 20      |

**Table 53 Command output**

| Field                  | Description                                                                                                                               |
|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| Authentication servers | Total number of RADIUS authentication servers.                                                                                            |
| Accounting servers     | Total number of RADIUS accounting servers.                                                                                                |
| IP                     | IP address of a RADIUS server.                                                                                                            |
| VPN                    | MPLS L3VPN instance to which the RADIUS server belongs.<br>This field displays <b>N/A</b> if no VPN instance is specified for the server. |
| Port                   | Service port number of the RADIUS server.                                                                                                 |
| Last 5 sec             | Total number of RADIUS authentication or accounting requests sent to the RADIUS server within the last 5 seconds.                         |
| History                | Total number of RADIUS authentication or accounting requests sent to the RADIUS server since the device starts up.                        |

## Related commands

```
reset radius server-load statistics
```

## reset radius server-load statistics

Use **reset radius server-load statistics** to clear history authentication and accounting load statistics for all RADIUS servers.

## Syntax

```
reset radius server-load statistics
```

## Views

User view

## Predefined user roles

network-admin

## Usage guidelines

This command does not clear authentication and accounting load statistics in the last 5 seconds.

## Examples

# Clear history authentication and accounting load statistics for all RADIUS servers.

```
<Sysname> reset radius server-load statistics
```

## Related commands

`display radius server-load statistics`

# New feature: 802.1X offline detection

## Configuring 802.1X offline detection

### About 802.1X offline detection

The 802.1X offline detection feature monitors the online status of 802.1X users. This feature uses an offline detect timer to set the interval that the device waits for traffic from a user before the device determines that the user is idle. If the device has not received traffic from a user before the timer expires, the device logs off that user and requests the accounting server to stop accounting for the user.

### Restrictions and guidelines

The 802.1X offline detection feature takes effect only on a port that performs MAC-based access control. If you change the port access mode to port-based, the 802.1X offline detection feature cannot take effect.

For this feature to operate as expected, do not set the offline detect timer to the same value as either of the following timers:

- Handshake timer (set by using the `dot1x timer handshake-period` command).
- Periodic reauthentication timer (set by using the `dot1x timer reauth-period` command).

### Procedure

1. Enter system view.  
`system-view`
2. Set the 802.1X offline detect timer.  
`dot1x timer offline-detect offline-detect-value`  
By default, the 802.1X offline detect timer is 300 seconds.
3. Enter interface view.  
`interface interface-type interface-number`
4. Enable 802.1X offline detection.  
`dot1x offline-detect enable`  
By default, 802.1X offline detection is disabled.

## Command reference

### dot1x offline-detect enable

Use `dot1x offline-detect enable` to enable 802.1X offline detection on a port.

Use `undo dot1x offline-detect enable` to disable 802.1X offline detection.

### Syntax

`dot1x offline-detect enable`  
`undo dot1x offline-detect enable`

### Default

802.1X offline detection is disabled on a port.

## Views

Layer 2 Ethernet interface view

## Predefined user roles

network-admin

## Usage guidelines

The 802.1X offline detection feature monitors the online status of 802.1X users. This feature uses an offline detect timer to set the interval that the device waits for traffic from a user before the device determines that the user is idle. If the device has not received traffic from a user before the timer expires, the device logs off that user and requests the accounting server to stop accounting for the user.

To have 802.1X offline detection take effect, you must configure the port to perform MAC-based access control. If you change the port access mode to port-based, the 802.1X offline detection feature cannot take effect.

To set the offline detect timer, use the **dot1x timer** command.

## Examples

# Disable 802.1X offline detection on Ten-GigabitEthernet 1/0/1.

```
<Sysname> system-view
[Sysname] interface ten-gigabitethernet 1/0/1
[Sysname-Ten-GigabitEthernet1/0/1] undo dot1x offline-detect enable
```

## Related commands

```
display dot1x
dot1x port-method
dot1x timer
```

## dot1x timer offline-detect

Use **dot1x timer offline-detect** to set the 802.1X offline detect timer.

Use **undo dot1x timer offline-detect** to restore the default.

## Syntax

```
dot1x timer offline-detect offline-detect-value
undo dot1x timer offline-detect
```

## Default

The 802.1X offline detect timer is 300 seconds.

## Views

System view

## Predefined user roles

network-admin

## Parameters

**offline-detect** *offline-detect-value*: Sets the offline detect timer in seconds. The value range is 60 to 2147483647.

## Usage guidelines

The offline detect timer sets the interval that the device waits for traffic from a user before the device determines that the user is idle. If the device has not received traffic from a user before the timer

expires, the device logs off that user and requests the accounting server to stop accounting for the user. This timer takes effect only when the 802.1X offline detection feature is enabled.

Change to the offline detect timer takes effect immediately on users.

## Examples

# Set the 802.1X offline detect timer to 150 seconds.

```
<Sysname> system-view
[Sysname] dot1x timer offline-detect 150
```

## Related commands

**display dot1x**

# New feature: 802.1X unauthenticated user aging

## Configuring 802.1X unauthenticated user aging

### About 802.1X unauthenticated user aging

802.1X unauthenticated user aging applies to users added to an 802.1X guest, critical, or Auth-Fail VLAN or VSI because they have not been authenticated or have failed authentication.

When a user in one of those VLANs or VSIs ages out, the device removes the user from the VLAN or VSI and deletes the MAC address entry for the user from the access port.

For users in one of those VLANs or VSIs on one port to be authenticated successfully and come online on another port, enable this feature. In any other scenarios, disable this feature as a best practice.

The 802.1X user aging mechanism on a port depends on its access control mode.

- If the port uses port-based access control, a user aging timer starts when the port is assigned to the critical or Auth-Fail VLAN or VSI. When the aging timer expires, the port is removed from the VLAN or VSI and all MAC address entries for users in the VLAN or VSI are also removed.
- If the port uses MAC-based access control, a user aging timer starts for each 802.1X user when they are assigned to the Auth-Fail, critical, or guest VLAN or VSI. When the aging timer for a user expires, the device removes that user from the VLAN or VSI.

The removed users will be unable to access any network resources until after another authentication is triggered.

### Restrictions and guidelines

As a best practice, disable 802.1X unauthenticated user aging on one port unless you want to have unauthenticated users on that port to be authenticated and come online on another port.

### Procedure

1. Enter system view.  
**system-view**
2. Set the user aging timer for a type of 802.1X VLAN or VSI.  
**dot1x timer user-aging { auth-fail-vlan | auth-fail-vsi | critical-vlan | critical-vsi | guest-vlan | guest-vsi } aging-time-value**  
The default setting is 1000 seconds for all aging timers.
3. Enter interface view.  
**interface interface-type interface-number**
4. Enable 802.1X unauthenticated user aging.

**dot1x unauthenticated-user aging enable**

- By default, 802.1X unauthenticated user aging is enabled.

## Command reference

### dot1x unauthenticated-user aging enable

Use **dot1x unauthenticated-user aging enable** to enable 802.1X unauthenticated user aging.

Use **undo dot1x unauthenticated-user aging enable** to disable 802.1X unauthenticated user aging.

#### Syntax

**dot1x unauthenticated-user aging enable**

**undo dot1x unauthenticated-user aging enable**

#### Default

User aging is enabled for 802.1X users that have not been authenticated or have not passed authentication.

#### Views

Layer 2 Ethernet interface view

#### Predefined user roles

network-admin

#### Usage guidelines

802.1X unauthenticated user aging applies to users added to 802.1X guest, critical, or Auth-Fail VLANs or VSIs because they have not been authenticated or have failed authentication.

When a user in one of those VLANs or VSIs ages out, the device removes the user from the VLAN or VSI and deletes the MAC address entry for the user from the access port.

For users in one of those VLANs or VSIs on one port to be authenticated successfully and come online on another port, enable this feature. In any other scenarios, disable this feature as a best practice.

The 802.1X user aging mechanism on a port depends on its access control mode.

- If the port uses port-based access control, a user aging timer starts when the port is assigned to the critical or Auth-Fail VLAN or VSI. When the aging timer expires, the port is removed from the VLAN or VSI and all MAC address entries for users in the VLAN or VSI are also removed.
- If the port uses MAC-based access control, a user aging timer starts for each 802.1X user when they are assigned to the Auth-Fail, critical, or guest VLAN or VSI. When the aging timer for a user expires, the device removes that user from the VLAN or VSI.

The removed users will be unable to access any network resources until after another authentication is triggered.

#### Examples

# Disable 802.1X user aging on Ten-GigabitEthernet 1/0/1.

```
<Sysname> system-view
```

```
[Sysname] interface ten-gigabitethernet 1/0/1
```

```
[Sysname-Ten-GigabitEthernet1/0/1] undo dot1x unauthenticated-user aging enable
```

#### Related commands

**dot1x timer**

## dot1x timer user-aging

Use **dot1x timer user-aging** to set the user aging timer for a type of 802.1X VLAN or VSI.

Use **undo dot1x timer user-aging** to restore the default user aging timer value for a type of 802.1X VLAN or VSI.

### Syntax

```
dot1x timer user-aging { auth-fail-vlan | auth-fail-vsi | critical-vlan
| critical-vsi | guest-vlan | guest-vsi } aging-time-value
undo dot1x timer user-aging { auth-fail-vlan | auth-fail-vsi |
critical-vlan | critical-vsi | guest-vlan | guest-vsi }
```

### Default

The user aging timers for all applicable types of 802.1X VLANs and VSIs are 1000 seconds.

### Views

System view

### Predefined user roles

network-admin

### Parameters

**auth-fail-vlan**: Specifies 802.1X Auth-Fail VLANs.

**auth-fail-vsi**: Specifies 802.1X Auth-Fail VSIs.

**critical-vlan**: Specifies 802.1X critical VLANs.

**critical-vsi**: Specifies 802.1X critical VSIs.

**guest-vlan**: Specifies 802.1X guest VLANs.

**guest-vsi**: Specifies 802.1X guest VSIs.

**aging-time-value**: Sets the user aging timer. The value range is 60 to 2147483647 seconds.

### Usage guidelines

If you enable 802.1X unauthenticated user aging, you can set a user aging timer for Auth-Fail, critical, or guest VLANs or VSIs. The user aging timer for a type of 802.1X VLAN or VSI determines how long a user can stay in that type of VLAN or VSI.

For more information about how user aging operates, see the usage guidelines for the **dot1x unauthenticated-user aging enable** command.

Do not set a user aging timer to an integral multiple of the username request timeout timer (the **dot1x timer tx-period** command). If you do so, the aging timer will not take effect.

A user aging timer change takes effect immediately on users.

### Examples

# Set the aging timer to 150 seconds for users in 802.1X critical VLANs.

```
<Sysname> system-view
```

```
[Sysname] dot1x timer user-aging critical-vlan 150
```

### Related commands

**display dot1x**

**dot1x unauthenticated-user aging enable**

# New feature: 802.1X EAP-TLS fragmentation for packets sent to the server

## Setting the maximum length of an EAP-TLS fragment sent to the server

### About 802.1X EAP-TLS fragmentation

When the device uses EAP-TLS authentication method in EAP relay mode, the RADIUS packets might exceed the maximum packet size supported by the RADIUS server. This situation typically occurs because long EAP-TLS messages are encapsulated in the EAP-Message attribute of the RADIUS packet sent to the RADIUS server.

To avoid authentication failures caused by oversized packets, fragment the EAP-TLS messages depending on the maximum RADIUS packet size supported by the remote RADIUS server.

For example, the maximum packet length allowed by the server is 1200 bytes and the length of a RADIUS packet (excluding the EAP-Message attribute) is 800 bytes. To make sure the maximum length of a RADIUS packet does not exceed 1200 bytes, you must set the maximum length of an EAP-TLS fragment to a value less than 400 bytes.

### Restrictions and guidelines

802.1X EAP-TLS fragmentation takes effect only when EAP relay mode is used.

### Procedure

1. Enter system view.  
**system-view**
2. Enable 802.1X EAP-TLS fragmentation and set the maximum EAP-TLS fragment size.  
**dot1x eap-tls-fragment to-server** *eap-tls-max-length*  
By default, EAP-TLS messages are not fragmented.

## Command reference

### dot1x eap-tls-fragment to-server

Use **dot1x eap-tls-fragment to-server** to enable 802.1X EAP-TLS fragmentation and set the maximum EAP-TLS fragment size.

Use **undo dot1x eap-tls-fragment to-server** to restore the default.

### Syntax

```
dot1x eap-tls-fragment to-server eap-tls-max-length
undo dot1x eap-tls-fragment to-server
```

### Default

EAP-TLS messages are not fragmented.

### Views

System view

### Predefined user roles

network-admin

## Parameters

*eap-tls-max-length*: Sets the maximum EAP-TLS fragment size in bytes. The value range is 100 to 1500.

## Usage guidelines

802.1X EAP-TLS fragmentation takes effect only when EAP relay mode is used.

When the device uses EAP-TLS authentication method in EAP relay mode, the RADIUS packets might exceed the maximum packet size supported by the RADIUS server. This situation typically occurs because long EAP-TLS messages are encapsulated in the EAP-Message attribute of the RADIUS packet sent to the RADIUS server.

To avoid authentication failures caused by oversized packets, fragment the EAP-TLS messages depending on the maximum RADIUS packet size supported by the remote RADIUS server.

For example, the maximum packet length allowed by the server is 1200 bytes and the length of a RADIUS packet (excluding the EAP-Message attribute) is 800 bytes. To make sure the maximum length of a RADIUS packet does not exceed 1200 bytes, you must set the maximum length of an EAP-TLS fragment to a value less than 400 bytes.

## Examples

```
Set the maximum EAP-TLS fragment size to 400 bytes.
```

```
<Sysname> system-view
```

```
[Sysname] dot1x eap-tls-fragment to-server 400
```

## Related commands

```
dot1x authentication-method
```

# New feature: User aging for unauthenticated MAC authentication users

## Configuring user aging for unauthenticated MAC authentication users

### About user aging for unauthenticated MAC authentication users

User aging for unauthenticated MAC authentication users applies to users added to a MAC authentication guest or critical VLAN or VSI because they have not been authenticated or have failed authentication.

When a user in one of those VLANs or VSIs ages out, the device removes the user from the VLAN or VSI and deletes the MAC address entry for the user from the access port.

For users in one of those VLANs or VSIs on one port to be authenticated successfully and come online on another port, enable this feature. In any other scenarios, disable this feature as a best practice.

### Restrictions and guidelines

As a best practice, disable user aging for unauthenticated MAC authentication users on one port unless you want to have the unauthenticated users on that port to be authenticated and come online on another port.

### Procedure

1. Enter system view.  
**system-view**

2. Set the user aging timer for a type of MAC authentication VLAN or VSI.  
**mac-authentication timer user-aging { critical-vlan | critical-vsi | guest-vlan | guest-vsi } aging-time-value**  
The default setting is 1000 seconds for all aging timers.
3. Enter interface view.  
**interface** *interface-type interface-number*
4. Enable user aging for unauthenticated MAC authentication users.  
**mac-authentication unauthenticated-user aging enable**  
By default, user aging is enabled for unauthenticated MAC authentication users.

## Command reference

### mac-authentication unauthenticated-user aging enable

Use **mac-authentication unauthenticated-user aging enable** to enable user aging for unauthenticated MAC authentication users.

Use **undo mac-authentication unauthenticated-user aging enable** to disable user aging for unauthenticated MAC authentication users.

#### Syntax

```
mac-authentication unauthenticated-user aging enable
undo mac-authentication unauthenticated-user aging enable
```

#### Default

User aging is enabled for unauthenticated MAC authentication users.

#### Views

Layer 2 Ethernet interface view

#### Predefined user roles

network-admin

#### Usage guidelines

User aging for unauthenticated MAC authentication users applies to users added to a MAC authentication guest or critical VLAN or VSI because they have not been authenticated or have failed authentication.

When a user in one of those VLANs or VSIs ages out, the device removes the user from the VLAN or VSI and deletes the MAC address entry for the user from the access port.

For users in one of those VLANs or VSIs on one port to be authenticated successfully and come online on another port, enable this feature. In any other scenarios, disable this feature as a best practice.

#### Examples

```
Disable user aging for unauthenticated MAC authentication users on Ten-GigabitEthernet 1/0/1.
<Sysname> system-view
[Sysname] interface ten-gigabitethernet 1/0/1
[Sysname-Ten-GigabitEthernet1/0/1] undo mac-authentication unauthenticated-user aging enable
```

#### Related commands

```
mac-authentication timer
```

## mac-authentication timer user-aging

Use **mac-authentication timer user-aging** to set the aging timer for a type of MAC authentication VLAN or VSI.

Use **undo mac-authentication timer user-aging** to restore the default user aging timer for a type of MAC authentication VLAN or VSI.

### Syntax

```
mac-authentication timer user-aging { critical-vlan | critical-vsi |
guest-vlan | guest-vsi } aging-time-value

undo mac-authentication timer user-aging { critical-vlan | critical-vsi
| guest-vlan | guest-vsi }
```

### Default

The user aging timer is 1000 seconds for all applicable types of MAC authentication VLANs and VSIs.

### Views

System view

### Predefined user roles

network-admin

### Parameters

**critical-vlan**: Specifies MAC authentication critical VLANs.

**critical-vsi**: Specifies MAC authentication critical VSIs.

**guest-vlan**: Specifies MAC authentication guest VLANs.

**guest-vsi**: Specifies MAC authentication guest VSIs.

*aging-time-value*: Sets the user aging timer. The value range is 60 to 2147483647 seconds.

### Usage guidelines

If you enable user aging for unauthenticated MAC authentication user, you can set a user aging timer for MAC authentication critical or guest VLANs or VSIs. The user aging timer for a type of MAC authentication VLAN or VSI determines how long a user can stay in that type of VLAN or VSI.

For more information about how user aging operates, see the usage guidelines for the **mac-authentication unauthenticated-user aging enable** command.

Do not set the user aging timer for users in MAC authentication guest VLANs or VSIs to an integral multiple of the authentication interval for them. If you do so, the aging timer will not take effect. The authentication interval for MAC authentication users in a guest VLAN or VSI is configurable with the **mac-authentication guest-vlan auth-period** or **mac-authentication guest-vsi auth-period** command, respectively.

A user aging timer change takes effect immediately on users.

### Examples

# Set the aging timer to 150 seconds for users in MAC authentication critical VLANs.

```
<Sysname> system-view
```

```
[Sysname] ac-authentication timer user-aging critical-vlan 150
```

### Related commands

**display mac-authentication**

**mac-authentication unauthenticated-user aging enable**

# New feature: MAC authentication method

## Specifying a MAC authentication method

### About MAC authentication methods

RADIUS-based MAC authentication supports the following authentication methods:

- **PAP**—Transports usernames and passwords in plain text. The authentication method applies to scenarios that do not require high security.
- **CHAP**—Transports usernames in plain text and passwords in encrypted form over the network. CHAP is more secure than PAP.

### Restrictions and guidelines

The device must use the same authentication method as the RADIUS server.

### Procedure

1. Enter system view.  
`system-view`
2. Specify an authentication method for MAC authentication.  
`mac-authentication authentication-method { chap | pap }`  
By default, the device uses PAP for MAC authentication.

## Command reference

### mac-authentication authentication-method

Use `mac-authentication authentication-method` to specify an authentication method for MAC authentication.

Use `undo mac-authentication authentication-method` to restore the default.

### Syntax

```
mac-authentication authentication-method { chap | pap }
undo mac-authentication authentication-method
```

### Default

The device uses PAP for MAC authentication.

### Views

System view

### Predefined user roles

network-admin

### Parameters

**chap**: Configures the access device to use the Challenge Handshake Authentication Protocol (CHAP) to communicate with the RADIUS server.

**pap**: Configures the access device to use the Password Authentication Protocol (PAP) to communicate with the RADIUS server.

### Usage guidelines

RADIUS-based MAC authentication supports the following authentication methods:

- **PAP**—Transports usernames and passwords in plain text. The authentication method applies to scenarios that do not require high security.
- **CHAP**—Transports usernames in plain text and passwords in encrypted form over the network. CHAP is more secure than PAP.

## Examples

```
Configure the device to use CHAP for MAC authentication.
<Sysname> system-view
[Sysname] mac-authentication authentication-method chap
```

## Related commands

```
display mac-authentication
```

# New feature: User-specific MAC authentication offline detection

## Configuring offline detection for a specific MAC authentication user

### About user-specific offline detection in MAC authentication

In addition to port-based MAC authentication offline detection, you can configure offline detection parameters on a per-user basis, as follows:

- Set an offline detect timer specific to a user and control whether to use the ARP snooping or ND snooping table to determine the offline state of the user.
  - If the ARP snooping or ND snooping table is used, the device searches the ARP snooping or ND snooping table before it checks for traffic from the user within the detection interval. If a matching ARP snooping or ND snooping entry is found, the device resets the offline detect timer and the user stays online. If the offline detect timer expires because the device has not found a matching snooping entry for the user or received traffic from the user, the device disconnects the user.
  - If the ARP or ND snooping table is not used, the device disconnects the user if it has not received traffic from that user before the offline detect timer expires.

When disconnecting the user, the device also notifies the RADIUS server (if any) to stop user accounting.

- Skip offline detection for the user. You can choose this option if the user is a dumb terminal. A dumb terminal might fail to come online again after it is logged off by the offline detection feature.

The device uses the offline detection settings for a user in the following sequence:

1. User-specific offline detection settings.
2. Offline detection settings assigned to the user by the RADIUS server. The settings include the offline detect timer, use of the ARP or ND snooping table in offline detection, and whether to ignore offline detection.
3. Port-based offline detection settings.

## Restrictions and guidelines

For the user-specific offline detection feature to take effect on a user, make sure the MAC authentication offline detection feature is enabled on the user's access port.

The user-specific offline detection settings take effect on the online users immediately after they are configured.

## Procedure

1. Enter system view.  
**system-view**
2. (Optional.) Configure MAC authentication offline detection for a user.  
**mac-authentication offline-detect mac-address *mac-address* { ignore | timer *offline-detect-value* [ check-arp-or-nd-snooping ] }**  
By default, offline detection settings configured on access ports take effect and the offline detect timer set in system view is used.

## Command reference

### mac-authentication offline-detect mac-address

Use **mac-authentication offline-detect mac-address** to configure MAC authentication offline detection for a MAC authentication user.

Use **undo mac-authentication offline-detect mac-address** to restore the default.

#### Syntax

```
mac-authentication offline-detect mac-address mac-address { ignore | timer offline-detect-value [check-arp-or-nd-snooping] }
undo mac-authentication offline-detect mac-address mac-address
```

#### Default

The offline detection settings configured on access ports take effect and the offline detect timer set in system view is used.

#### Views

System view

#### Predefined user roles

network-admin

#### Parameters

***mac-address***: Specifies a MAC address in the format of H-H-H, excluding multicast, all-zero, and all-F MAC addresses.

**ignore**: Skips offline detection for the specified user.

**timer *offline-detect-value***: Specifies the offline detect timer for the specified user. The value range is 60 to 2147483647 seconds.

**check-arp-or-nd-snooping**: Uses the ARP snooping or ND snooping table in offline detection to determine the offline state of the user.

#### Usage guidelines

Use this command to set offline detection parameters specific to a MAC authentication user. To have this command take effect, you must make sure MAC authentication offline detection is enabled on the user's access port. The user-specific offline detection settings take effect on the online users immediately after they are configured.

Use this command as follows:

- Set an offline detect timer specific to a user and control whether to use the ARP snooping or ND snooping table to determine the offline state of the user.
  - If the ARP snooping or ND snooping table is used, the device searches the ARP snooping or ND snooping table before it checks for traffic from the user within the detection interval. If

a matching ARP snooping or ND snooping entry is found, the device resets the offline detect timer and the user stays online. If the offline detect timer expires because the device has not found a matching snooping entry for the user or received traffic from the user, the device disconnects the user.

- If the ARP or ND snooping table is not used, the device disconnects the user if it has not received traffic from that user before the offline detect timer expires.

When disconnecting the user, the device also notifies the RADIUS server (if any) to stop user accounting.

- Skip offline detection for the user. You can choose this option if the user is a dumb terminal. A dumb terminal might fail to come online again after it is logged off by the offline detection feature.

The device uses the offline detection settings for a user in the following sequence:

1. User-specific offline detection settings.
2. Offline detection settings assigned to the user by the RADIUS server. The settings include the offline detect timer, use of the ARP or ND snooping table in offline detection, and whether to ignore offline detection.
3. Port-based offline detection settings.
4. Examples

# Disable MAC authentication offline detection for the MAC authentication user with MAC address 000a-eb29-7511.

```
<Sysname> system-view
```

```
[Sysname] mac-authentication offline-detect mac-address 000a-eb29-7511 ignore
```

# Enable MAC authentication offline detection for the MAC authentication user with MAC address 000a-eb29-7511, and set the offline detect timer to 24 hours.

```
<Sysname> system-view
```

```
[Sysname] mac-authentication offline-detect mac-address 000a-eb29-7511 timer 86400
```

## Related commands

```
display mac-authentication connection
```

```
mac-authentication offline-detect enable
```

```
mac-authentication timer (system view)
```

# New feature: NAS-Port-Type attribute

## Configuring the NAS-Port-Type attribute

### About the NAS-Port-Type attribute

The NAS-Port-Type attribute in a RADIUS request represents the type of a user's access interface.

The access device might not be able to correctly obtain the type of users' access interfaces when multiple network devices exist between the access device and the portal client. For the access device to send the correct access interface type to the RADIUS server, perform this task to configure the NAS-Port-Type attribute.

### Restrictions and guidelines

This configuration takes effect only on portal users that newly come online.

### Procedure

1. Enter system view.

**system-view**

2. Enter interface view.

**interface** *interface-type interface-number*

3. Configure the NAS-Port-Type attribute carried in outgoing RADIUS requests on the interface.

```
portal nas-port-type { 802.11 | adsl-cap | adsl-dmt | async | cable
| ethernet | g.3-fax | hdlc | idsl | isdn-async-v110 | isdn-async-v120
| isdn-sync | piafs | sds1 | sync | virtual | wireless-other | x.25
| x.75 | xds1 }
```

By default, the NAS-Port-Type carried in outgoing RADIUS requests is Ethernet (attribute value 15).

## Command reference

### portal nas-port-type

Use **portal nas-port-type** to configure the NAS-Port-Type attribute carried in outgoing RADIUS requests.

Use **undo portal nas-port-type** to restore the default.

#### Syntax

```
portal nas-port-type { 802.11 | adsl-cap | adsl-dmt | async | cable |
ethernet | g.3-fax | hdlc | idsl | isdn-async-v110 | isdn-async-v120 |
isdn-sync | piafs | sds1 | sync | virtual | wireless-other | x.25 | x.75
| xds1 }
```

```
undo portal nas-port-type
```

#### Default

The NAS-Port-Type carried in outgoing RADIUS requests is Ethernet (attribute value 15).

#### Views

Interface view

#### Predefined user roles

network-admin

#### Parameters

**802.11**: Specifies the NAS port type as Wireless-IEEE 802.11 (attribute value 19).

**adsl-cap**: Specifies the NAS port type as ADSL-CAP (attribute value 12).

**adsl-dmt**: Specifies the NAS port type as ADSL-DMT (attribute value 13).

**async**: Specifies the NAS port type as Async (attribute value 0).

**cable**: Specifies the NAS port type as Cable (attribute value 17).

**ethernet**: Specifies the NAS port type as Ethernet (attribute value 15).

**g.3-fax**: Specifies the NAS port type as G.3 Fax (attribute value 10).

**hdlc**: Specifies the NAS port type as HDLC Clear Channel (attribute value 7).

**ids1**: Specifies the NAS port type as IDSL (attribute value 14).

**isdn-async-v110**: Specifies the NAS port type as ISDN Async V.110 (attribute value 4).

**isdn-async-v120**: Specifies the NAS port type as ISDN Async V.120 (attribute value 3).

**isdn-sync**: Specifies the NAS port type as ISDN Sync (attribute value 2).  
**piafs**: Specifies the NAS port type as PIAFS (attribute value 6).  
**sds1**: Specifies the NAS port type as SDSL (attribute value 11).  
**sync**: Specifies the NAS port type as Sync (attribute value 1).  
**virtual**: Specifies the NAS port type as Virtual (attribute value 5).  
**wireless-other**: Specifies the NAS port type as Wireless-Other (attribute value 18).  
**x.25**: Specifies the NAS port type as X.25 (attribute value 8).  
**x.75**: Specifies the NAS port type as X.75 (attribute value 9).  
**xds1**: Specifies the NAS port type as xDSL (attribute value 16).

## Examples

```
Configure the NAS-Port-Type carried in outgoing RADIUS requests as SDSL on VLAN-interface 20.
<Sysname> system-view
[Sysname] interface vlan-interface 20
[Sysname-Vlan-interface20] portal nas-port-type sds1
```

# New feature: Checking the issuing of category-2 portal filtering rules

## Checking the issuing of category-2 portal filtering rules

### About checking the issuing of category-2 portal filtering rules

Category-2 portal filtering rules permit authenticated users to access authorized network resources. By default, the device allows an authenticated user to come online as long as a member device has issued a category-2 portal filtering rule for the user. Users coming online from global interfaces might fail to access network resources because some member ports might not have category-2 rules for the users. To resolve this issue, enable the device to check the issuing of category-2 portal filtering rules. Then, the device allows users to come online only when all member devices have issued category-2 portal filtering rules for the users.

As a best practice, perform this task when portal authentication is enabled on a global interface.

### Procedure

1. Enter system view.  
**system-view**
2. Enable the device to check the issuing of category-2 portal filtering rules.  
**portal user-rule assign-check enable**  
By default, the device does not check the issuing of category-2 portal filtering rules.

## Command reference

### portal user-rule assign-check enable

Use **portal user-rule assign-check enable** to enable the device to check the issuing of category-2 portal filtering rules.

Use **undo portal user-rule assign-check enable** to disable the device from checking the issuing of category-2 portal filtering rules.

### Syntax

```
portal user-rule assign-check enable
undo portal user-rule assign-check enable
```

### Default

The device does not check the issuing of category-2 portal filtering rules.

### Views

System view

### Predefined user roles

network-admin

### Usage guidelines

To display category-2 portal filtering rules, use the **display portal rule dynamic** command.

### Examples

```
Enable the device to check the issuing of category-2 portal filtering rules.
<Sysname> system-view
[Sysname] portal user-rule assign-check enable
```

### Related commands

```
display portal rule dynamic
```

## New feature: Configuring user authentication timeout timer for global password control

### Configuring user authentication timeout timer for global password control

#### About configuring user authentication timeout timer for global password control

After global password control is enabled, you can configure the authentication timeout timer for the Telnet and terminal users that are attempting to log in to the device. If a user does not pass authentication within the timeout time, the connection is torn down.

#### Procedure

1. Enter system view.  
**system-view**
2. Configure the user authentication timeout timer.  
**password-control authentication-timeout** *timeout*  
By default, the user authentication timeout timer is 600 seconds.  
This feature only takes effect for Telnet and terminal users.

## Command reference

### password-control authentication-timeout

Use `password-control authentication-timeout` to configure the user authentication timeout timer.

Use `undo password-control authentication-timeout` to restore the default.

#### Syntax

```
password-control authentication-timeout timeout
undo password-control authentication-timeout
```

#### Default

The user authentication timeout timer is 600 seconds.

#### Views

System view

#### Predefined user roles

network-admin

#### Usage guidelines

This feature only takes effect for Telnet and terminal users. If a user does not pass authentication within the timeout time, the connection is torn down.

#### Examples

```
Set the user authentication timeout timer to 40 seconds.
<Sysname> system-view
[Sysname] password-control authentication-timeout 40
```

#### Related commands

```
display password-control
```

## New feature: Configuring password control for local network access users

### Setting global password control features for network access users

#### About global password control features for network access users

From this release, support of password control for network access users was added.

- You can enable password control in system view for all network access users by specifying the **network-class** keyword.
- You can configure the minimum password length, password complexity checking policy, and password composition policy in user group view and local user view of network access users.

#### Restrictions and guidelines

You can configure the following password control features for network access users:

- Minimum password length.
- Password complexity checking policy.

- Password composition policy.
- Minimum password update interval.
- Maximum number of history password records.

The minimum password length, password complexity checking policy, and password composition policy can be configured in system view, user group view, and local user view. The password settings with a smaller application scope have higher priority. For local users, password settings configured in local user view have the highest priority, and global settings in system view have the lowest priority.

## Procedure

1. Enter system view.  
**system-view**
2. Enable global password control for network access users.  
**password-control enable network-class**  
By default, global password control is disabled for network access users.
3. Enable a specific password control feature.  
**password-control { composition | history | length } enable**  
By default, all password control features are enabled.
4. Configure password settings.
  - o Set the minimum password length.  
In non-FIPS mode:  
**password-control length length**  
The default setting is 10 characters.  
In FIPS mode:  
**password-control length length**  
The default setting is 15 characters.
  - o Configure the password composition policy.  
In non-FIPS mode:  
**password-control composition type-number type-number**  
**[ type-length type-length ]**  
By default, a password must contain a minimum of one character type and a minimum of one character for each type.  
In FIPS mode:  
**password-control composition type-number type-number**  
**[ type-length type-length ]**  
By default, a password must contain a minimum of four character types and a minimum of one character for each type.
  - o Configure the password complexity checking policy.  
**password-control complexity { same-character | user-name } check**  
By default, the system does not perform password complexity checking.
  - o Set the maximum number of history password records for each user.  
**password-control history max-record-number**  
The default setting is 4.
  - o Set the minimum password update interval.  
**password-control update-interval interval**  
The default setting is 24 hours.

## Setting password control parameters for a single network access user

1. Enter system view.  
**system-view**
2. Create a network access user and enter network access user view.  
**local-user** *user-name* **class network**  
For information about configuring a local user, see "Configuring AAA."
3. Configure the minimum password length for the local user.  
**password-control length** *length*  
By default, the setting equals that for the user group to which the local user belongs. If no minimum password length is configured for the user group, the global setting applies to the local user.
4. Configure the password composition policy for the local user.  
**password-control composition type-number** *type-number* [ **type-length** *type-length* ]  
By default, the settings equal those for the user group to which the local user belongs. If no password composition policy is configured for the user group, the global settings apply to the local user.
5. Configure the password complexity checking policy for the local user.  
**password-control complexity** { **same-character** | **user-name** } **check**  
By default, the settings equal those for the user group to which the local user belongs. If no password complexity checking policy is configured for the user group, the global settings apply to the local user.

## Command reference

### Modified command: display password-control

#### Syntax

```
display password-control
```

#### Views

Any view

#### Change description

Before modification: You can display only the password control configuration for device management users.

After modification: The following password control settings for network access users were added to the command output:

- Password control configuration.
- User authentication timeout timer.
- Time when the password history records of the device management or network access users were deleted.

#### Examples

```
Display the password control configuration.
```

```
<Sysname> display password-control
```

```
Global password control configurations:
```

|                                       |                                          |
|---------------------------------------|------------------------------------------|
| Password control:                     | Enabled (device management users)        |
|                                       | Enabled (network access users)           |
| Password aging:                       | Enabled (90 days)                        |
| Password length:                      | Enabled (10 characters)                  |
| Password composition:                 | Enabled (1 types, 1 characters per type) |
| Password history:                     | Enabled (max history records:4)          |
| Early notice on password expiration:  | 7 days                                   |
| User authentication timeout:          | 600 seconds                              |
| Maximum login attempts:               | 3                                        |
| Action for exceeding login attempts:  | Lock user for 1 minutes                  |
| Password history was last reset:      | 0 days ago (device management users)     |
|                                       | 0 days ago (network access users)        |
| Minimum interval between two updates: | 24 hours                                 |
| User account idle time:               | 90 days                                  |
| Logins with aged password:            | 3 times in 30 days                       |
| Password complexity:                  | Disabled (username checking)             |
|                                       | Disabled (repeated characters checking)  |

**Table 54 Command output**

| Field                                | Description                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Password control                     | Whether the password control feature is enabled for device management users and network access users.                                                                                                                                                                                                                                                            |
| Password aging                       | Whether password expiration is enabled and, if enabled, the aging time.                                                                                                                                                                                                                                                                                          |
| Password length                      | Whether the minimum password length restriction feature is enabled and, if enabled, the setting.                                                                                                                                                                                                                                                                 |
| Password composition                 | Whether the password composition restriction feature is enabled and, if enabled, the settings.                                                                                                                                                                                                                                                                   |
| Password history                     | Whether the password history feature is enabled and, if enabled, the setting.                                                                                                                                                                                                                                                                                    |
| Early notice on password expiration  | Number of days during which the user is notified of the pending password expiration.                                                                                                                                                                                                                                                                             |
| User authentication timeout          | User authentication timeout timer in seconds.                                                                                                                                                                                                                                                                                                                    |
| Maximum login attempts               | Allowed maximum number of consecutive failed login attempts for FTP and VTY users.                                                                                                                                                                                                                                                                               |
| Action for exceeding login attempts  | Action to be taken after a user fails to log in after the specified number of attempts.                                                                                                                                                                                                                                                                          |
| Password history was last reset      | Last time when the password history records of the device management or network access users were deleted.                                                                                                                                                                                                                                                       |
| Minimum interval between two updates | Minimum password update interval.                                                                                                                                                                                                                                                                                                                                |
| Logins with aged password            | Number of times and maximum number of days a user can log in using an expired password.                                                                                                                                                                                                                                                                          |
| Password complexity                  | Whether the following password complexity checking is enabled: <ul style="list-style-type: none"> <li><b>username checking</b>—Checks whether a password contains the username or the reverse of the username.</li> <li><b>repeated characters checking</b>—Checks whether a password contains any character that appears consecutively three or more</li> </ul> |

| Field | Description |
|-------|-------------|
|       | times.      |

## Modified feature: password-control complexity

### Syntax

```
password-control complexity { same-character | user-name } check
undo password-control complexity { same-character | user-name } check
```

### Views

System view

User group view

Local user view

### Change description

Before modification:

- In system view or user group view, this command configures the password complexity checking policy only for device management users.
- This command is only supported in local user view of device management users.

After modification:

- In system view or user group view, you can use this command to configure the password complexity checking policy for device management users and network access users.
- This command is supported in local user view of both device management users and network access users.

## Modified feature: password-control composition

### Syntax

```
password-control composition type-number type-number [type-length
type-length]
undo password-control composition
```

### Views

System view

User group view

Local user view

### Change description

Before modification:

- In system view or user group view, this command applies only to the passwords of device management users.
- This command is only supported in local user view of device management users.

After modification:

- In system view or user group view, this command applies to the passwords of device management users and network access users.
- This command is supported in local user view of both device management users and network access users.

## Modified feature: password-control enable

### Old syntax

```
password-control enable
undo password-control enable
```

### New syntax

```
password-control enable [network-class]
undo password-control enable [network-class]
```

### Views

System view

### Change description

Before modification: You can enable the password control feature in system view only for device management users.

After modification: You can enable the password control feature in system view for network access users by specifying the **network-class** keyword.

## Modified feature: password-control history

### Syntax

```
password-control history max-record-number
undo password-control history
```

### Views

System view

### Change description

Before modification: This command in system view only applies to passwords of device management users.

After modification: This command in system view applies to passwords of both device management users and network access users.

## Modified feature: password-control length

### Syntax

```
password-control length length
undo password-control length
```

### Views

System view

User group view

Local user view

### Change description

Before modification:

- This command in system view or user group view only applies to passwords of device management users.
- This command is supported in local user view of device management users.

After modification:

- This command in system view or user group view applies to passwords of device management users and network access users.
- This command is supported in local user view of both device management users and network access users.

## Modified feature: password-control update-interval

### Syntax

```
password-control update-interval interval
undo password-control update-interval
```

### Views

System view

### Change description

Before modification: This command in system view only applies to passwords of device management users.

After modification: This command in system view applies to passwords of both device management users and network access users.

## Modified feature: reset password-control history-record

### Old syntax

```
reset password-control history-record [super [role role-name] |
user-name user-name]
```

### New syntax

```
reset password-control history-record [super [role role-name] |
user-name user-name | network-class [user-name user-name]]
```

### Views

User view

### Change description

Before modification: This command deletes history records of super passwords and history passwords of device management users.

After modification: You can also use this command to delete history password records of network access users by specifying the **network-class** keyword.

## New feature: Configuring password change at first login

### Configuring password change at first login

#### About configuring password change at first login

By default, if the global password control feature is enabled, users must change the password at first login before they can access the system by default. If it is not necessary for users to change the password at first login, disable the password change at first login feature.

#### Procedure

1. Enter system view.

**system-view**

2. Disable password change at first login.

**undo password-control change-password first-login enable**

By default, the password change at first login feature is enabled.

In FIPS mode, the password change at first login feature cannot be disabled.

## Command reference

### password-control change-password first-login enable

Use **password-control change-password first-login enable** to enable the password change at first login feature.

Use **undo password-control change-password first-login enable** to restore the default.

#### Syntax

**password-control change-password first-login enable**

**undo password-control change-password first-login enable**

#### Default

The password change at first login feature is enabled.

#### Views

System view

#### Predefined user roles

network-admin

#### Usage guidelines

For the password change at first login feature to take effect, make sure the global password control is enabled.

In FIPS mode, the password change at first login feature cannot be disabled.

#### Examples

# Enable password change at first login.

```
<Sysname> system-view
```

```
[Sysname] password-control change-password first-login enable
```

#### Related command

**display password-control**

**password-control enable**

# New feature: Enabling the SSL server to send the complete certificate chain during SSL negotiation

## Enabling the SSL server to send the complete certificate chain during SSL negotiation

About enabling the SSL server to send the complete certificate chain during SSL negotiation

To ensure successful SSL session establishment, configure this feature if the client does not have the complete certificate chain to verify the server certificate.

### Procedure

1. Enter system view.  
**system-view**
2. Create an SSL server policy and enter SSL server policy view.  
**ssl server-policy** *policy-name*
3. Enable the SSL server to send the complete certificate chain to the client during SSL negotiation.  
**certificate-chain-sending enable**

By default, during SSL negotiation, the SSL server sends the server certificate rather than the complete certificate chain to the client.

## Command reference

### certificate-chain-sending enable

Use **certificate-chain-sending enable** to enable the SSL server to send the complete certificate chain to the client during SSL negotiation.

Use **undo certificate-chain-sending enable** to restore the default.

### Syntax

**certificate-chain-sending enable**  
**undo certificate-chain-sending enable**

### Default

During SSL negotiation, the SSL server sends the server certificate rather than the complete certificate chain to the client.

### Views

SSL server policy view

### Predefined user roles

network-admin

### Usage guidelines

This feature causes additional overheads in the SSL negotiation process. Enable it only when the SSL client does not have the complete certificate chain to verify the server certificate.

### Examples

# Enable the SSL server to send the complete certificate chain to the client during SSL negotiation.

```
<Sysname> system-view
[Sysname] ssl server-policy policy1
[Sysname-ssl-server-policy-policy1] certificate-chain-sending enable
```

## New feature: Support of RRPP for the link-up delay timer

### Configuring the link-up delay timer

#### About configuring the link-up delay timer

This feature prevents frequent switchover of RRPP traffic forwarding paths caused by unstable RRPP port states. This feature behaves differently depending on whether you specify the **distribute** keyword in the **linkup-delay-timer** command.

- If you do not specify the **distribute** keyword, the master node starts the link-up delay timer when a faulty port comes up and the master node receives Hello packets from the secondary port.
  - If the master node can still receive Hello packets from the secondary port after the link-up delay timer expires, the master node performs the following operations:
    - Changes the RRPP ring state from Disconnect to Health.
    - Switches the traffic from the secondary port to the primary port.
  - If the master node cannot receive Hello packets from the secondary port after the Fail timer expires and before the link-up delay timer expires, the master node performs the following operations:
    - Stops the link-up delay timer.
    - Keeps the RRPP ring in Disconnect state.
- If you specify the **distribute** keyword, all nodes in the RRPP domain can learn the value of the link-up delay timer through Hello packets. When the faulty port comes up, the master node performs the following operations:
  - The hosting RRPP node shuts down the faulty port (the faulty port cannot send or receive any packets).
  - Starts the link-up delay timer.

If the port does not become faulty after the link-up delay timer expires, the hosting RRPP node sets the port state to up. The master node can receive Hello packets from its secondary port again. Then, the master node changes the RRPP ring state from Disconnect to Health and switches the traffic from the secondary port to the primary port.

If the port becomes faulty again before the link-up delay timer expires, the hosting RRPP node blocks the port and stops the link-up delay timer.

#### Restrictions and guidelines

If the **distribute** keyword is not specified, the link-up delay timer value cannot be greater than the Fail timer value minus twice the Hello timer value.

#### Procedure

1. Enter system view.  
**system-view**
2. Enter RRPP domain view.  
**rrpp domain** *domain-id*
3. Configure the link-up delay timer.  
**linkup-delay-timer** *delay-time* [ **distribute** ]

By default, the link-up delay timer is 0 seconds, and the **distribute** keyword is not specified.

## Command reference

### linkup-delay-timer

Use **linkup-delay-timer** to configure the link-up delay timer.

Use **undo linkup-delay-timer** to restore the default.

#### Syntax

```
linkup-delay-timer delay-time [distribute]
undo linkup-delay-timer
```

#### Default

The link-up delay timer is 0 seconds, and the **distribute** keyword is not specified.

#### Views

RRPP domain view

#### Predefined user roles

network-admin

#### Parameters

*delay-time*: Specifies the link-up delay timer value in the range of 0 to 30 seconds.

**distribute**: Enables all nodes in the RRPP domain to learn the link-up delay timer value.

#### Usage guidelines

This command prevents frequent switchover of RRPP traffic forwarding paths caused by unstable RRPP port states.

This command can be configured on all nodes in the RRPP ring, but it can take effect only on the master node.

If the **distribute** keyword is not specified, the link-up delay timer value cannot be greater than the Fail timer value minus twice the Hello timer value.

#### Examples

# Set the link-up delay timer to 10 seconds for RRPP domain 1.

```
<Sysname> system-view
[Sysname] rrpp domain 1
[Sysname-rrpp-domain1] linkup-delay-timer 10
```

#### Related commands

**timer**

## New feature: gRPC

### About gRPC

gRPC is an open source remote procedure call (RPC) system initially developed at Google. It uses HTTP 2.0 for transport and provides network device configuration and management methods that support multiple programming languages.

## gRPC protocol stack layers

Table 55 describes the gRPC protocol stack layers.

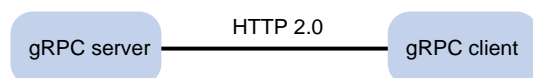
**Table 55 gRPC protocol stack layers**

| Layer                           | Description                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------|
| Content layer                   | Defines the data of the service module.<br>Two peers must notify each other of the data models that they are using. |
| Protocol buffers encoding layer | Encodes data by using the protocol buffers code format.                                                             |
| gRPC layer                      | Defines the protocol interaction format for remote procedure calls.                                                 |
| HTTP 2.0 layer                  | Carries gRPC.                                                                                                       |
| TCP layer                       | Provides connection-oriented reliable serial data links.                                                            |

## Network architecture

As shown in Figure 16, the gRPC network uses the client/server model. It uses HTTP 2.0 for packet transport.

**Figure 16 gRPC network architecture**



The gRPC network uses the following mechanism:

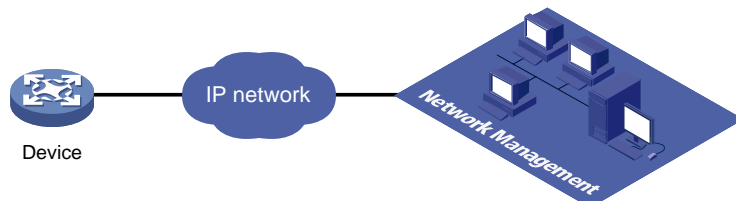
2. The gRPC server listens to connection requests from clients at the gRPC service port.
3. A user runs the gRPC client application to log in to the gRPC server, and uses methods provided in the .proto file to send requests.
4. The gRPC server responds to requests from the gRPC client.

The device can act as the gRPC server or client.

## Telemetry technology based on gRPC

Telemetry is a remote data collection technology for monitoring device performance and operating status. Hewlett Packard Enterprise telemetry technology uses gRPC to push data from the device to the collectors on the NMSs. As shown in Figure 17, after a gRPC connection is established between the device and NMSs, the NMSs can subscribe to data of modules on the device.

**Figure 17 Telemetry technology based on gRPC**



## Telemetry modes

The device supports the following telemetry modes:

- **Dial-in mode**—The device acts as a gRPC server and the collectors act as gRPC clients. A collector initiates a gRPC connection to the device to subscribe to device data.  
Dial-in mode applies to small networks where collectors need to deploy configurations to devices.
- **Dial-out mode**—The device acts as a gRPC client and the collectors act as gRPC servers. The device initiates a gRPC connection to the collectors and pushes subscribed device data to the collectors.  
Dial-out mode applies to larger networks where devices need to push device data to collectors.

## Protocols

RFC 7540, *Hypertext Transfer Protocol version 2 (HTTP/2)*

## FIPS compliance

The device supports the FIPS mode that complies with NIST FIPS 140-2 requirements. Support for features, commands, and parameters might differ in FIPS mode and non-FIPS mode. For more information about FIPS mode, see *Security Configuration Guide*.

gRPC is not supported in FIPS mode.

## Configuring the gRPC dial-in mode

### gRPC dial-in mode configuration tasks at a glance

To configure the gRPC dial-in mode, perform the following tasks:

1. [Configuring the gRPC service](#)
2. [Configuring a gRPC user](#)

### Configuring the gRPC service

1. Enter system view.  
**system-view**
2. (Optional.) Set the gRPC service port number.  
**grpc port** *port-number*  
By default, the gRPC service port number is 50051.
3. Enable the gRPC service.  
**grpc enable**  
By default, the gRPC service is disabled.
4. (Optional.) Set the gRPC session idle timeout timer.  
**grpc idle-timeout** *minutes*  
By default, the gRPC session idle timeout timer is 5 minutes.

### Configuring a gRPC user

#### About gRPC users

For gRPC clients to establish gRPC sessions with the device, you must configure local users for the gRPC clients and authorize the local users to use the HTTPS service.

## Procedure

1. Enter system view.  
**system-view**
2. Add a local user with the device management right.  
**local-user** *user-name* [ **class** **manage** ]
3. Configure a password for the user.  
**password** [ { **hash** | **simple** } *password* ]  
By default, no password is configured for a local user. A non-password-protected user can pass authentication after providing the correct username and passing attribute checks.
4. (Optional.) Assign a user role to the user.  
**authorization-attribute** **user-role** *user-role*  
By default, a local user is assigned the network-operator role.
5. Authorize the user to use the HTTPS service.  
**service-type** **https**  
By default, no service types are authorized to a local user.

For more information about the **local-user**, **password**, **authorization-attribute**, and **service-type** commands, see AAA configuration in *Security Command Reference*.

## Configuring the gRPC dial-out mode

### gRPC dial-out mode configuration tasks at a glance

To configure telemetry, perform the following tasks:

1. [Enabling the gRPC service](#)
2. [Configuring sensors](#)
3. [Configuring collectors](#)
4. [Configuring a subscription](#)

### Enabling the gRPC service

1. Enter system view.  
**system-view**
2. Enable the gRPC service.  
**grpc enable**  
By default, the gRPC service is disabled.

## Configuring sensors

### About sensors

The device uses sensors to sample data. A sensor path indicates a data source.

Supported data sampling types include:

- **Event-triggered sampling**—Sensors in a sensor group sample data when certain events occur. For sensor paths of this data sampling type, see *NETCONF XML API Event Reference* for the module.

- **Periodic sampling**—Sensors in a sensor group sample data at intervals. For sensor paths of this data sampling type, see the NETCONF XML API references for the module except for *NETCONF XML API Event Reference*.

## Procedure

1. Enter system view.  
**system-view**
2. Enter telemetry view.  
**telemetry**
3. Create a sensor group and enter sensor group view.  
**sensor-group** *group-name*
4. Specify a sensor path.  
**sensor path** *path*  
To specify multiple sensor paths, execute this command multiple times.

## Configuring collectors

### About collectors

Collectors are used to receive sampled data from network devices. For the device to communicate with collectors, you must create a destination group and add collectors to the destination group.

### Restrictions and guidelines

As a best practice, configure a maximum of five destination groups, and specify a maximum of five collectors in a destination group. If you configure too many destination groups or specify too many collectors in a destination group, system performance might degrade.

## Procedure

1. Enter system view.  
**system-view**
2. Enter telemetry view.  
**telemetry**
3. Create a destination group and enter destination group view.  
**destination-group** *group-name*
4. Specify a collector.  
IPv4:  
**ipv4-address** *ipv4-address* [ **port** *port-number* ] [ **vpn-instance** *vpn-instance-name* ]  
IPv6:  
**ipv6-address** *ipv6-address* [ **port** *port-number* ] [ **vpn-instance** *vpn-instance-name* ]  
To specify multiple collectors, execute this command multiple times.

## Configuring a subscription

### About configuring a subscription

A subscription binds sensor groups to destination groups. Then, the device pushes data from the specified sensors to the collectors.

## Procedure

1. Enter system view.  
**system-view**
2. Enter telemetry view.  
**telemetry**
3. Create a subscription and enter subscription view.  
**subscription** *subscription-name*
4. (Optional.) Specify the source IP address for packets sent to collectors.  
**source-address** { *ipv4-address* | **interface** *interface-type interface-number* | **ipv6** *ipv6-address* }  
  
By default, the device uses the primary IPv4 address of the output interface for the route to the collectors as the source address.
5. Specify a sensor group.  
**sensor-group** *group-name* [ **sample-interval** *interval* ]  
  
Specify the **sample-interval** *interval* option for periodic sensor paths and only for periodic sensor paths.
  - If you specify the option for event-triggered sensor paths, the sensor paths do not take effect.
  - If you do not specify the option for periodic sensor paths, the device does not sample or push data.
6. Specify a destination group.  
**destination-group** *group-name*

## Display and maintenance commands for gRPC

Execute **display** commands in any view.

| Task                      | Command             |
|---------------------------|---------------------|
| Display gRPC information. | <b>display grpc</b> |

## gRPC configuration examples

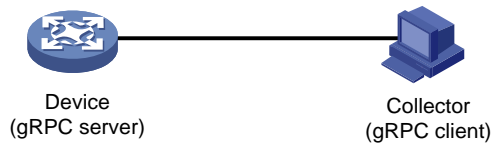
These configuration examples describe only CLI configuration tasks on the device. The collectors need to run an extra application. For information about collector-side application development, see "[Developing the collector-side application](#)."

### Example: Configuring the gRPC dial-in mode

#### Network configuration

As shown in [Figure 18](#), configure the gRPC dial-in mode on the device so the device acts as the gRPC server and the gRPC client can subscribe to LLDP events on the device.

**Figure 18 Network diagram**



## Procedure

1. Assign IP addresses to interfaces on the gRPC server and client and configure routes. Make sure the server and client can reach each other.
2. Configure the device as the gRPC server:
  - # Enable the gRPC service.

```
<Device> system-view
[Device] grpc enable
```
  - # Create a local user named **test**. Set the password to **test**, and assign the HTTPS service to the user.

```
[Device] local-user test
[Device-luser-manage-test] password simple test
[Device-luser-manage-test] service-type https
[Device-luser-manage-test] quit
```
3. Configure the gRPC client.
  - a. Prepare a PC and install the gRPC environment on the PC. For more information, see the user guide for the gRPC environment.
  - b. Obtain the Hewlett Packard Enterprise proto definition file and uses the protocol buffers compiler to generate code of a specific language, for example, Java, Python, C/C++, or Go.
  - c. Create a client application to call the generated code.
  - d. Start the application to log in to the gRPC server.

## Verifying the configuration

When an LLDP event occurs on the gRPC server, verify that the gRPC client receives the event.

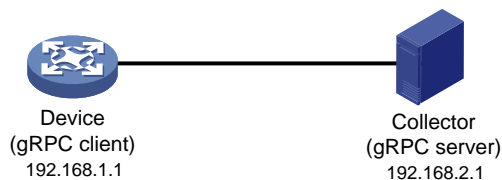
## Example: Configuring the gRPC dial-out mode

### Network configuration

As shown in [Figure 19](#), the device is connected to a collector. The collector uses port 50050.

Configure gRPC dial-out mode on the device so the device pushes the device capability information of its interface module to the collector at 10-second intervals.

**Figure 19 Network diagram**



## Procedure

# Configure IP addresses as required so the device and the collector can reach each other. (Details not shown.)

```

Enable the gRPC service.
<Device> system-view
[Device] grpc enable

Create a sensor group named test, and add sensor path ifmgr/devicecapabilities/.
[Device] telemetry
[Device-telemetry] sensor-group test
[Device-telemetry-sensor-group-test] sensor path ifmgr/devicecapabilities/
[Device-telemetry-sensor-group-test] quit

Create a destination group named collector1. Specify a collector that uses IPv4 address
192.168.2.1 and port number 50050.
[Device-telemetry] destination-group collector1
[Device-telemetry-destination-group-collector1] ipv4-address 192.168.2.1 port 50050
[Device-telemetry-destination-group-collector1] quit

Configure a subscription named A to bind sensor group test with destination group collector1. Set
the sampling interval to 10 seconds.
[Device-telemetry] subscription A
[Device-telemetry-subscription-A] sensor-group test sample-interval 10
[Device-telemetry-subscription-A] destination-group collector1
[Device-telemetry-subscription-A] quit

```

## Verifying the configuration

```

Verify that the collector receives the device capability information of the interface module from the
device at 10-second intervals. (Details not shown.)

```

## gRPC dial-in mode commands

### display grpc

Use **display grpc** to display gRPC dial-in mode information.

#### Syntax

```
display grpc
```

#### Views

Any view

#### Predefined user roles

```

network-admin
network-operator

```

#### Examples

```

Display gRPC dial-in mode information.
<Sysname> display grpc
gRPC status : enabled.
gRPC port : 50051
gRPC idle-timeout : 3 minutes
Session count: 1.
 Session ID: 1
 User name: test
 Login time:2018-01-05 06:46:43 Idle time : 2 mins 56 s

```

```

Client IP address : 169.254.100.170:40810
Received RPCs : 0 Received error RPCs : 0
Received subscription: 0 Output notifications: 0

```

**Table 56 Command output**

| Field                 | Description                                                                                                                                                                      |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| gRPC status           | Status of the gRPC service: <ul style="list-style-type: none"> <li><b>enabled</b>—The gRPC service is enabled.</li> <li><b>disabled</b>—The gRPC service is disabled.</li> </ul> |
| gRPC idle-timeout     | Setting for the gRPC session idle timeout timer.                                                                                                                                 |
| Session count         | Number of gRPC sessions.                                                                                                                                                         |
| Idle time             | Duration in which the session idle timeout timer will expire. If the value of this field is 0, gRPC sessions will never be timed out.                                            |
| Received error RPCs   | Number of received erroneous gRPC requests.                                                                                                                                      |
| Received subscription | Number of received gRPC subscription requests.                                                                                                                                   |

## grpc enable

Use **grpc enable** to enable the gRPC service.

Use **undo grpc enable** to disable the gRPC service.

### Syntax

```
grpc enable
```

```
undo grpc enable
```

### Default

The gRPC service is disabled.

### Views

System view

### Predefined user roles

network-admin

### Examples

```
Enable the gRPC service.
```

```
<Sysname> system
```

```
[Sysname] grpc enable
```

## grpc idle-timeout

Use **grpc idle-timeout** to set the gRPC session idle timeout timer.

Use **undo grpc idle-timeout** to restore the default.

### Syntax

```
grpc idle-timeout minutes
```

```
undo grpc idle-timeout
```

## Default

The gRPC session idle timeout timer is 5 minutes.

## Views

System view

## Predefined user roles

network-admin

## Parameters

*minutes*: Specifies the gRPC session idle timeout timer in minutes, in the range of 0 to 30. To disable gRPC sessions from being timed out, set it to 0.

## Usage guidelines

If no gRPC packet exchanges occur on the session between a gRPC and the server before the idle timeout timer expires, the device closes the session.

## Examples

```
Set the gRPC session idle timeout timer to 6 minutes.
<Sysname> system
[Sysname] grpc idle-timeout 6
```

## grpc port

Use **grpc port** to specify the gRPC service port number.

Use **undo grpc port** to restore the default.

## Syntax

```
grpc port port-number
undo grpc port
```

## Default

The gRPC service port number is 50051.

## Views

System view

## Predefined user roles

network-admin

## Parameters

*port-number*: Specifies the gRPC service port number, in the range of 1 to 65535.

## Usage guidelines

You can configure this command only when the gRPC service is disabled.

## Examples

```
Set the gRPC service port number to 50052.
<Sysname> system
[Sysname] grpc port 50052
```

## Related commands

```
grpc enable
```

## gRPC dial-out mode commands

### destination-group (subscription view)

Use **destination-group** to specify a destination group for a subscription.

Use **undo destination-group** to remove a destination group from a subscription.

#### Syntax

**destination-group** *group-name*

**undo destination-group** *group-name*

#### Default

A subscription does not have a destination group.

#### Views

Subscription view

#### Predefined user roles

network-admin

#### Parameters

*group-name*: Specifies a destination group by its name, a case-sensitive string of 1 to 31 characters.

#### Usage guidelines

A subscription binds sensor groups to destination groups. Then, the device pushes data from the specified sensors to the collectors.

The specified destination group must have been created by using the **destination-group** command in telemetry view.

#### Examples

# Specify destination group **collector1** for subscription **A**.

```
<Sysname> system-view
```

```
[Sysname] telemetry
```

```
[Sysname-telemetry] subscription A
```

```
[Sysname-telemetry-subscription-A] destination-group collector1
```

#### Related commands

**destination-group** (telemetry view)

### destination-group (telemetry view)

Use **destination-group** to create a destination group and enter its view, or enter the view of an existing destination group.

Use **undo destination-group** to delete a destination group.

#### Syntax

**destination-group** *group-name*

**undo destination-group** *group-name*

#### Default

No destination groups exist.

## Views

Telemetry view

## Predefined user roles

network-admin

## Parameters

*group-name*: Specifies the destination group name, a case-sensitive string of 1 to 31 characters.

## Usage guidelines

As a best practice, configure a maximum of five destination groups, and specify a maximum of five collectors in a destination group. If you configure too many destination groups or specify too many collectors in a destination group, system performance might degrade.

## Examples

```
Create a destination group named collector1.
<Sysname> system-view
[Sysname] telemetry
[Sysname-telemetry] destination-group collector1
[Sysname-telemetry-destination-group-collector1]
```

## ipv4-address

Use **ipv4-address** to add an IPv4 collector to a destination group.

Use **undo ipv4-address** to remove an IPv4 collector from a destination group.

## Syntax

```
ipv4-address ipv4-address [port port-number] [vpn-instance vpn-instance-name]
undo ipv4-address ipv4-address [port port-number] [vpn-instance vpn-instance-name]
```

## Default

A destination group does not have IPv4 collectors.

## Views

Destination group view

## Predefined user roles

network-admin

## Parameters

*ipv4-address*: Specifies the IPv4 address of the collector.

**port** *port-number*: Specifies the listening port of the collector, in the range of 1 to 65535. The default is 50051.

**vpn-instance** *vpn-instance-name*: Specifies the VPN instance to which the collector belongs. The *vpn-instance-name* argument represents the VPN instance name, a case-sensitive string of 1 to 31 characters. If the collector belongs to the public network, do not specify this option.

## Usage guidelines

To add multiple collectors to a destination group, execute this command multiple times.

One collector must have a different address, port, or VPN instance than the other collectors.

## Examples

# Add a collector that uses IPv4 address 192.168.21.21 and the default port number to destination group **collector1**.

```
<Sysname> system-view
[Sysname] telemetry
[Sysname-telemetry] destination-group collector1
[Sysname-telemetry-destination-group-collector1] ipv4-address 192.168.21.21
```

## Related commands

**destination-group** (telemetry view)

## ipv6-address

Use **ipv6-address** to add an IPv6 collector to a destination group.

Use **undo ipv6-address** to remove an IPv6 collector from a destination group.

## Syntax

```
ipv6-address ipv6-address [port port-number] [vpn-instance vpn-instance-name]
undo ipv6-address ipv6-address [port port-number] [vpn-instance vpn-instance-name]
```

## Default

A destination group does not have IPv6 collectors.

## Views

Destination group view

## Predefined user roles

network-admin

## Parameters

*ipv6-address*: Specifies the IPv6 address of the collector.

**port** *port-number*: Specifies the listening port of the collector, in the range of 1 to 65535. The default is 50051.

**vpn-instance** *vpn-instance-name*: Specifies the VPN instance to which the collector belongs. The *vpn-instance-name* argument represents the VPN instance name, a case-sensitive string of 1 to 31 characters. If the collector belongs to the public network, do not specify this option.

## Usage guidelines

To add multiple collectors to a destination group, execute this command multiple times.

One collector must have a different address, port, or VPN instance than the other collectors.

## Examples

# Add a collector that uses IPv6 address 1::1 and the default port number to destination group **collector1**.

```
<Sysname> system-view
[Sysname] telemetry
[Sysname-telemetry] destination-group collector1
[Sysname-telemetry-destination-group-collector1] ipv6-address 1::1
```

## Related commands

**destination-group** (telemetry view)

## sensor path

Use **sensor path** to configure a sensor path.

Use **undo sensor path** to delete a sensor path.

## Syntax

**sensor path** *path*

**undo sensor path** *path*

## Default

No sensor paths exist.

## Views

Sensor group view

## Predefined user roles

network-admin

## Parameters

*path*: Specifies a data path. For information about the available paths, enter a question mark (?) in the position of this argument.

## Usage guidelines

To configure multiple sensor paths, execute this command multiple times.

## Examples

```
Configure sensor path ifmgr/devicecapabilities/ for sensor group test.
<Sysname> system-view
[Sysname] telemetry
[Sysname-telemetry] sensor-group test
[Sysname-telemetry-sensor-group-test] sensor path ifmgr/devicecapabilities/
```

## Related commands

**sensor-group** (telemetry view)

## sensor-group (subscription view)

Use **sensor-group** to specify a sensor group for a subscription.

Use **undo sensor-group** to remove a sensor group from a subscription.

## Syntax

**sensor-group** *group-name* [ **sample-interval** *interval* ]

**undo sensor-group** *group-name*

## Default

A subscription does not have a sensor group.

## Views

Subscription view

## Predefined user roles

network-admin

## Parameters

*group-name*: Specifies a sensor group by its name, a case-sensitive string of 1 to 31 characters.

**sample-interval** *interval*: Specifies the data sampling interval in seconds. The value range is 1 to 86400.

## Usage guidelines

Specify the **sample-interval** *interval* option for periodic sensor paths and only for periodic sensor paths.

- If you specify the option for event-triggered sensor paths, the sensor paths do not take effect.
- If you do not specify the option for periodic sensor paths, the device does not sample or push data.

The specified sensor group must have been created by using the **sensor-group** command in telemetry view.

## Examples

# Specify sensor group **test** for subscription **A**. Set the data sampling interval to 10 seconds.

```
<Sysname> system-view
```

```
[Sysname] telemetry
```

```
[Device-telemetry] subscription A
```

```
[Device-telemetry-subscription-A] sensor-group test sample-interval 10
```

## Related commands

**sensor path**

**sensor-group** (telemetry view)

## sensor-group (telemetry view)

Use **sensor-group** to create a sensor group and enter its view, or enter the view of an existing sensor group.

Use **undo sensor-group** to delete a sensor group.

## Syntax

**sensor-group** *group-name*

**undo sensor-group** *group-name*

## Default

No sensor groups exist.

## Views

Telemetry view

## Predefined user roles

network-admin

## Parameters

*group-name*: Specifies the sensor group name, a case-sensitive string of 1 to 31 characters.

## Examples

# Create a sensor group named **test**.

```

<Sysname> system-view
[Sysname] telemetry
[Sysname-telemetry] sensor-group test
[Sysname-telemetry-sensor-group-test]

```

## source-address

Use **source-address** to specify the source IP address for packets sent to collectors.

Use **undo source-address** to restore the default.

### Syntax

```

source-address { ipv4-address | interface interface-type
interface-number | ipv6 ipv6-address }
undo source-address

```

### Default

The device uses the primary IPv4 address of the output interface for the route to the collectors as the source address.

### Views

Subscription view

### Predefined user roles

network-admin

### Parameters

*ipv4-address*: Specifies an IPv4 address.

**interface** *interface-type interface-number*: Specifies an interface by its type and number. In the current software version, you must specify a loopback interface. The device will use the interface's primary IPv4 address as the source address. If the interface does not have a primary IPv4 address, the device uses the primary IPv4 address of the output interface for the route to the collectors.

### Usage guidelines

If you execute this command multiple times, the most recent configuration takes effect.

### Examples

# Specify the source IPv4 address of 169.254.1.1 for packets sent to collectors.

```

<Sysname> system-view
[Sysname] telemetry
[Sysname-telemetry] subscription A
[Sysname-telemetry-subscription-A] source-address ip 169.254.1.1

```

## subscription

Use **subscription** to create a subscription and enter its view, or enter the view of an existing subscription.

Use **undo sensor-group** to delete a subscription.

### Syntax

```

subscription subscription-name
undo subscription subscription-name

```

## Default

No subscription groups exist.

## Views

Telemetry view

## Predefined user roles

network-admin

## Parameters

*subscription-name*: Specifies the subscription name, a case-sensitive string of 1 to 31 characters.

## Examples

```
Configure a subscription named A.
<Sysname> system-view
[Sysname] telemetry
[Sysname-telemetry] subscription A
[Sysname-telemetry-subscription-A]
```

## Related commands

**destination-group** (subscription view)

**sensor-group** (subscription view)

## telemetry

Use **telemetry** to enter telemetry view.

## Syntax

**telemetry**

## Views

System view

## Predefined user roles

network-admin

## Usage guidelines

In telemetry view, you can configure telemetry parameters.

## Examples

```
Enter telemetry view.
<Sysname> system-view
[Sysname] telemetry
[Sysname-telemetry]
```

# New feature: Protocol buffers

## Protocol buffers code format

Interface description language developed by Google. It is used to define data structures and generates code based on multiple programming languages. Compared with XML, protocol buffers is simpler and quicker in parsing.

Protocol buffers are a method of serializing structured data. The protocol buffers code is similar to the corresponding XML code and JSON code. The protocol buffers code is binary and provides higher performance.

Table 57 compares a protocol buffers code format example and the corresponding JSON code format example.

**Table 57 Protocol buffers and JSON code format examples**

| Protocol buffers code format example                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Corresponding JSON code format example                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <pre>{ 1:"HPE" 2:"HPE" 3:"HPE Simware" 4:"Syslog/LogBuffer" 5:"notification": {   "Syslog": {     "LogBuffer": {       "BufferSize": 512,       "BufferSizeLimit": 1024,       "DroppedLogsCount": 0,       "LogsCount": 100,       "LogsCountPerSeverity": {         "Alert": 0,         "Critical": 1,         "Debug": 0,         "Emergency": 0,         "Error": 3,         "Informational": 80,         "Notice": 15,         "Warning": 1       },       "OverwrittenLogsCount": 0,       "State": "enable"     }   },   "Timestamp": "1527206160022" }</pre> | <pre>{   "producerName": "HPE",   "deviceName": "HPE",   "deviceModel": "HPE Simware",   "sensorPath": "Syslog/LogBuffer",   "jsonData": {     "notification": {       "Syslog": {         "LogBuffer": {           "BufferSize": 512,           "BufferSizeLimit": 1024,           "DroppedLogsCount": 0,           "LogsCount": 100,           "LogsCountPerSeverity": {             "Alert": 0,             "Critical": 1,             "Debug": 0,             "Emergency": 0,             "Error": 3,             "Informational": 80,             "Notice": 15,             "Warning": 1           },           "OverwrittenLogsCount": 0,           "State": "enable"         }       },       "Timestamp": "1527206160022"     }   } }</pre> |

## Proto definition files

You can define data structures in a proto definition file. Then, you can compile the file with utility `protoc` to generate code in a programming language such as Java and C++. Using the generated code, you can develop an application for a collector to communicate with the device.

Hewlett Packard Enterprise provides proto definition files for both dial-in mode and dial-out mode.

## Proto definition files in dial-in mode

### Public proto definition files

The **grpc\_service.proto** file defines the public RPC methods in dial-in mode, for example, login method and logout method.

The following are the contents of the **grpc\_service.proto** file:

```
syntax = "proto2";
package grpc_service;

message GetJsonReply { // Reply to the Get method
 required string result = 1;
}

message SubscribeReply { // Subscription result
 required string result = 1;
}

message ConfigReply { // Configuration result
 required string result = 1;
}

message ReportEvent { // Subscribed event
 required string token_id = 1; // Login token_id
 required string stream_name = 2; // Event stream name
 required string event_name = 3; // Event name
 required string json_text = 4; // Subscription result, a JSON string
}

message GetReportRequest{ // Obtains the event subscription result
 required string token_id = 1; // Returns the token_id upon a successful login
}

message LoginRequest { // Login request parameters
 required string user_name = 1; // Username
 required string password = 2; // Password
}

message LoginReply { // Reply to a login request
 required string token_id = 1; // Returns the token_id upon a successful login
}

message LogoutRequest { // Logout parameter
 required string token_id = 1; // token_id
}

message LogoutReply { // Reply to a logout request
 required string result = 1; // Logout result
}

message SubscribeRequest { // Event stream name
 required string stream_name = 1;
```

```

}
service GrpcService { // gRPC methods
 rpc Login (LoginRequest) returns (LoginReply) {} // Login method
 rpc Logout (LogoutRequest) returns (LogoutReply) {} // Logout method
 rpc SubscribeByStreamName (SubscribeRequest) returns (SubscribeReply) {} // Event
subscription method
 rpc GetEventReport (GetReportRequest) returns (stream ReportEvent) {} // Method for
obtaining the subscribed event
}

```

## Proto definition files for service modules

The dial-in mode supports proto definition files for the following service modules: Device, Ifmgr, IPFW, LLDP, and Syslog.

The following are the contents of the **Device.proto** file, which defines the RPC methods for the Device module:

```

syntax = "proto2";
import "grpc_service.proto";
package device;

message DeviceBase { // Structure for obtaining basic device information
 optional string HostName = 1; // Device name
 optional string HostOid = 2; // sysoid
 optional uint32 MaxChassisNum = 3; //Maximum number of chassis
 optional uint32 MaxSlotNum = 4; // Maximum number of slots
 optional string HostDescription = 5; // Device description
}

message DevicePhysicalEntities { // Structure for obtaining physical entity information
of the device
 message Entity {
 optional uint32 PhysicalIndex = 1; // Entity index
 optional string VendorType = 2; // Vendor type
 optional uint32 EntityClass = 3; // Entity class
 optional string SoftwareRev = 4; // Software version
 optional string SerialNumber = 5; // Serial number
 optional string Model = 6; // Model
 }
 repeated Entity entity = 1;
}

service DeviceService { // RPC methods
 rpc GetJsonDeviceBase(DeviceBase) returns (grpc_service.GetJsonReply) {} // Method
for obtaining basic device information
 rpc GetJsonDevicePhysicalEntities(DevicePhysicalEntities) returns
(grpc_service.GetJsonReply) {} // Method for obtaining physical entity information of
the device
}

```

## Proto definition file in dial-out mode

The **grpc\_dialout.proto** file defines the public RPC methods in dial-out mode. The following are the contents of the file:

```

syntax = "proto2";

```

```

package grpc_dialout;

message DeviceInfo{ // Pushed device information
 required string producerName = 1; // Vendor name
 required string deviceName = 2; // Device name
 required string deviceModel = 3; // Device model
}

message DialoutMsg{ // Format of the pushed data
 required DeviceInfo deviceMsg = 1; // Device information described by DeviceInfo
 required string sensorPath = 2; // Sensor path, which corresponds to xpath in NETCONF
 required string jsonData = 3; // Sampled data, a JSON string
}

message DialoutResponse{ // Response from the collector. Reserved. The value is not
 processed.
 required string response = 1;
}

service gRPCDialout { // Data push method
 rpc Dialout(stream DialoutMsg) returns (DialoutResponse);
}

```

## Obtaining proto definition files

Contact the technical support.

## Example: Developing a gRPC collector-side application

Use a language (for example, C++) to develop a gRPC collector-side application on Linux to enable a collector to collect device data.

## Prerequisites

1. Obtain HPE proto definition files.
  - For dial-in mode, obtain the **grpc\_service.proto** file and proto definition files for service modules.
  - For dial-out mode, obtain the **grpc\_dialout.proto** file.
2. Obtain utility protoc from <https://github.com/google/protobuf/releases>.
3. Obtain the protobuf plug-in for C++ (protobuf-cpp) from <https://github.com/google/protobuf/releases>.

## Generating the C++ code for the proto definition files

### Dial-in mode

# Copy the required proto definition files to the current directory, for example, **grpc\_service.proto** and **BufferMonitor.proto**.

```
$ protoc --plugin=./grpc_cpp_plugin --grpc_out=. --cpp_out=. *.proto
```

### Dial-out mode

# Copy proto definition file **grpc\_dialout.proto** to the current directory.

```
$ protoc --plugin=./grpc_cpp_plugin --grpc_out=. --cpp_out=. *.proto
```

## Developing the collector-side application

### Dial-in mode

In dial-in mode, the application needs to provide the code to be run on the gRPC client.

The C++ code generated from the proto definition files already encapsulates the service classes, which are GrpcService and BufferMonitorService in this example. For the gRPC client to initiate RPC requests, you only need to call the RPC method in the application.

The application performs the following operations:

- Log in to obtain the token\_id.
- Prepare parameters for the RPC method, use the service classes generated from the proto definition files to call the RPC method, and resolve the returned result.
- Log out.

To develop the collector-side application in dial-in mode:

#### 1. Create a GrpcServiceTest class.

# In the class, use the GrpcService::Stub class generated from grpc\_service.proto. Implement login and logout with the Login and Logout methods generated from grpc\_service.proto.

```
class GrpcServiceTest
{
public:
 /* Constructor functions */
 GrpcServiceTest(std::shared_ptr<Channel> channel):
 GrpcServiceStub(GrpcService::NewStub(channel)) {}

 /* Member functions */
 int Login(const std::string& username, const std::string& password);
 void Logout();
 void listen();

 /* Member variable */
 std::string token;

private:
 std::unique_ptr<GrpcService::Stub> GrpcServiceStub; // Use the
 GrpcService::Stub class generated from grpc_service.proto.
};
```

#### 2. Customize the Login method.

# Call the Login method of the GrpcService::Stub class to allow a user who provides the correct the username and password to log in.

```
int GrpcServiceTest::Login(const std::string& username, const std::string& password)
{
 LoginRequest request; // Username and password.
 request.set_user_name(username);
 request.set_password(password);

 LoginReply reply;
 ClientContext context;

 // Call the Login method.
```

```

Status status = GrpcServiceStub->Login(&context, request, &reply);
if (status.ok())
{
 std::cout << "login ok!" << std::endl;
 std::cout <<"token id is :" << reply.token_id() << std::endl;
 token = reply.token_id(); // The login succeeds. The token is obtained.
 return 0;
}
else{
 std::cout << status.error_code() << ": " << status.error_message()
 << ". Login failed!" << std::endl;
 return -1;
}
}

```

3. Initiate an RPC request to the device. In this example, the application subscribes to interface packet drop events.

```

rpc SubscribePortQueueDropEvent(PortQueueDropEvent) returns
(grpc_service.SubscribeReply) {}

```

4. Create the BufMon\_GrpcClient class to encapsulate the RPC method.  
# Use the BufferMonitorService::Stub class generated from BufferMonitor.proto to call the RPC method.

```

class BufMon_GrpcClient
{
public:
 BufMon_GrpcClient(std::shared_ptr<Channel> channel):
mStub(BufferMonitorService::NewStub(channel))
 {
 }

 std::string BufMon_Sub_AllEvent(std::string token);
 std::string BufMon_Sub_BoardEvent(std::string token);
 std::string BufMon_Sub_PortOverrunEvent(std::string token);
 std::string BufMon_Sub_PortDropEvent(std::string token);

 /* Get entries */
 std::string BufMon_Sub_GetStatistics(std::string token);
 std::string BufMon_Sub_GetGlobalCfg(std::string token);
 std::string BufMon_Sub_GetBoardCfg(std::string token);
 std::string BufMon_Sub_GetNodeQueueCfg(std::string token);
 std::string BufMon_Sub_GetPortQueueCfg(std::string token);

private:
 std::unique_ptr<BufferMonitorService::Stub> mStub; // Use the class generated
from BufferMonitor.proto.
};

```

5. Use std::string BufMon\_Sub\_PortDropEvent(std::string token) to implement interface packet drop event subscription.

```

std::string BufMon_GrpcClient::BufMon_Sub_PortDropEvent(std::string token)
{

```

```

std::cout << "-----BufMon_Sub_PortDropEvent----- " << std::endl;

PortQueDropEvent stNodeEvent;
PortQueDropEvent_PortQueDrop* pstParam = stNodeEvent.add_portquedrop();

UINT uiIfIndex = 0;
UINT uiQueIdx = 0;
UINT uiAlarmType = 0;

std::cout<<"Please input interface queue info : ifIndex queIdx alarmtype " <<
std::endl;
cout<<"alarmtype : 1 for ingress; 2 for egress; 3 for port headroom"<<endl;

std::cin>>uiIfIndex>>uiQueIdx>>uiAlarmType; // Set the subscription parameters
and interface index.
pstParam->set_ifindex(uiIfIndex);
pstParam->set_queindex(uiQueIdx);
pstParam->set_alarmtype(uiAlarmType);

ClientContext context;

/* Token needs to be added to context */ // Set the token_id to be returned after
a successful login
std::string key = "token_id";
std::string value = token;
context.AddMetadata(key, value);

SubscribeReply reply;
Status status = mStub->SubscribePortQueDropEvent(&context,stNodeEvent,&reply);
// Call the RPC method.

return reply.result();
}

```

## 6. Use a loop to listen to event reports.

# Implement this method in the GrpcServiceTest class.

```

void GrpcServiceTest::listen()
{
 GetReportRequest reportRequest;
 ClientContext context;
 ReportEvent reportedEvent;

 /* Add the token to the request */
 reportRequest.set_token_id(token);

 std::unique_ptr< ClientReader< ReportEvent>>
reader(GrpcServiceStub->GetEventReport(&context, reportRequest)); // Use
GetEventReport (which is generated from grpc_service.proto) to obtain event
information.

```

```

 std::string streamName;
 std::string eventName;
 std::string jsonText;
 std::string token;

 JsonFormatTool jsonTool;

 std::cout << "Listen to server for Event" << std::endl;

 while(reader->Read(&reportedEvent)) // Read the received event report.
 {
 streamName = reportedEvent.stream_name();
 eventName = reportedEvent.event_name();
 jsonText = reportedEvent.json_text();
 token = reportedEvent.token_id();

 std::cout << "/******EVENT COME*****/"
<< std::endl;
 std::cout << "TOKEN: " << token << std::endl;
 std::cout << "StreamName: " << streamName << std::endl;
 std::cout << "EventName: " << eventName << std::endl;
 std::cout << "JsonText without format: " << std::endl << jsonText << std::endl;
 std::cout << std::endl;
 std::cout << "JsonText Formated: " << jsonTool.formatJson(jsonText) <<
std::endl;
 std::cout << std::endl;
 }

 Status status = reader->Finish();
 std::cout << "Status Message:" << status.error_message() << "ERROR code : " <<
status.error_code();
 } // Login and RPC request finished.

```

## 7. To log out, call the Logout method. (Details not shown.)

### Dial-out mode

In dial-out mode, the application needs to provide the gRPC server code so the collector can receive and resolve data obtained from the device.

The application performs the following operations:

- Inherit the automatically generated GRPCDialout::Service class, overload the automatically generated RPC Dialout service, and resolve the fields.
- Register the RPC service with the specified listening port.

To develop the collector-side application in dial-out mode:

#### 1. Inherit and overload RPC service Dialout.

# Create class DialoutTest and inherit GRPCDialout::Service.

```

class DialoutTest final : public GRPCDialout::Service { // Overload the automatically
generated abstract class.
 Status Dialout(ServerContext* context, ServerReader< DialoutMsg>* reader,
DialoutResponse* response) override; // Implement RPC method Dialout.
};

```

2. Register the DialoutTest service as a gRPC service and specify the listening port.

```
using grpc::Server;
using grpc::ServerBuilder;

std::string server_address("0.0.0.0:60057"); // Specify the address and port to
listen to.

DialoutTest dialout_test; // Define the object declared in step 1.
ServerBuilder builder;
builder.AddListeningPort(server_address, grpc::InsecureServerCredentials()); // Add
the listening port.
builder.RegisterService(&dialout_test); // Register the service.
std::unique_ptr<Server> server(builder.BuildAndStart()); // Start the service.
server->Wait();
```

3. Implement the Dialout method and data resolution.

```
Status DialoutTest::Dialout(ServerContext* context, ServerReader< DialoutMsg>*
reader, DialoutResponse* response)
{
 DialoutMsg msg;

 while(reader->Read(&msg))
 {
 const DeviceInfo &device_msg = msg.devicemsg();
 std::cout<< "Producer-Name: " << device_msg.producername() << std::endl;
 std::cout<< "Device-Name: " << device_msg.devicename() << std::endl;
 std::cout<< "Device-Model: " << device_msg.devicemodel() << std::endl;
 std::cout<<"Sensor-Path: " << msg.sensorpath()<<std::endl;
 std::cout<<"Json-Data: " << msg.jsondata()<<std::endl;
 std::cout<<std::endl;
 }
 response->set_response("test");

 return Status::OK;
}
```

4. After obtaining the DialoutMsg object (generated from the proto definition file) through the Read method, you can call the method to obtain the field values.

## New feature: Specifying the NTP time-offset thresholds for log and trap outputs

### Specifying the NTP time-offset thresholds for log and trap outputs

#### About NTP time-offset thresholds for log and trap outputs

When the time offset between an NTP client and its reference source exceeds 128 ms, the NTP client synchronizes its time to its reference source.

#### Procedure

1. Enter system view.  
**system-view**
2. Specify the NTP time-offset thresholds for log and trap outputs.

```
ntp-service time-offset-threshold { log log-threshold | trap
trap-threshold } *
```

By default, no NTP time-offset thresholds are set for log and trap outputs.

## Command reference

### ntp-service time-offset-threshold

Use **ntp-service time-offset-threshold** to specify the NTP time-offset thresholds for log and trap outputs.

Use **undo ntp-service time-offset-threshold** to restore the default.

#### Syntax

```
ntp-service time-offset-threshold { log log-threshold | trap
trap-threshold } *
undo ntp-service time-offset-threshold
```

#### Default

No NTP time-offset thresholds are set for log and trap outputs.

#### Views

System view

#### Predefined user roles

network-admin

#### Parameters

**log log-threshold**: Specifies the NTP time-offset threshold for log output. The value range for the *log-threshold* argument is 128 to 60000, in milliseconds.

**trap trap-threshold**: Specifies the NTP time-offset threshold for trap output. The value range for the *trap-threshold* argument is 128 to 60000, in milliseconds.

#### Usage guidelines

When the time offset between an NTP client and its reference source exceeds 128 ms, the NTP client synchronizes its time to its reference source.

#### Examples

```
Specify the NTP time-offset thresholds for log and trap outputs as 500 ms and 600 ms,
respectively.
```

```
<Sysname> system-view
```

```
[Sysname] ntp-service time-offset-threshold log 500 trap 600
```

## New feature: Specifying the SNTP time-offset thresholds for log and trap outputs

### Specifying the SNTP time-offset thresholds for log and trap outputs

#### About SNTP time-offset thresholds for log and trap outputs

When the time offset between an SNTP client and its reference source exceeds 128 ms, the SNTP client synchronizes its time to its reference source.

## Procedure

1. Enter system view.  
**system-view**
2. Specify the SNTP time-offset thresholds for log and trap outputs.  
**sntp time-offset-threshold { log *log-threshold* | trap *trap-threshold* }**  
\*

By default, no SNTP time-offset thresholds are set for log and trap outputs.

## Command reference

### sntp time-offset-threshold

Use **sntp time-offset-threshold** to specify the SNTP time-offset thresholds for log and trap outputs.

Use **undo sntp time-offset-threshold** to restore the default.

#### Syntax

```
sntp time-offset-threshold { log log-threshold | trap trap-threshold } *
undo sntp time-offset-threshold
```

#### Default

No SNTP time-offset thresholds are set for log and trap outputs.

#### Views

System view

#### Predefined user roles

network-admin

#### Parameters

**log *log-threshold***: Specifies the SNTP time-offset threshold for log output. The value range for the *log-threshold* argument is 128 to 60000, in milliseconds.

**trap *trap-threshold***: Specifies the SNTP time-offset threshold for trap output. The value range for the *trap-threshold* argument is 128 to 60000, in milliseconds.

#### Usage guidelines

When the time offset between an SNTP client and its reference source exceeds 128 ms, the SNTP client synchronizes its time to its reference source.

#### Examples

# Specify the SNTP time-offset thresholds for log and trap outputs as 500 ms and 600 ms, respectively.

```
<Sysname> system-view
```

```
[Sysname] sntp time-offset-threshold log 500 trap 600
```

# New feature: Ignoring the permit flag added by OpenFlow

## Ignoring the permit flag added by OpenFlow

### About ignoring the permit flag added by OpenFlow

Packets matching an OpenFlow flow entry with the output action as Normal are permitted to pass through and marked with a permit flag. Packets with a permit flag are not dropped by a discard action in a QoS policy (for example, a CAR discard action).

This feature allows the device to drop packets with a permit flag by using a discard action in a QoS policy.

### Procedure

1. Enter system view.  
**system-view**
2. Ignore the permit flag added by OpenFlow.  
**openflow permit-flag ignore**  
By default, the permit flag added by OpenFlow is not ignored.

## Command reference

### openflow permit-flag ignore

Use **openflow permit-flag ignore** to ignore the permit flag added by OpenFlow.

Use **undo openflow permit-flag ignore** to restore the default.

#### Syntax

```
openflow permit-flag ignore
undo openflow permit-flag ignore
```

#### Default

The permit flag added by OpenFlow is not ignored.

#### Views

System view

#### Predefined user roles

network-admin

#### Usage guidelines

Packets matching an OpenFlow flow entry with the output action as Normal are permitted to pass through and marked with a permit flag. Packets with a permit flag are not dropped by a discard action in a QoS policy (for example, a CAR discard action).

This command allows the device to drop packets with a permit flag by using a discard action in a QoS policy.

#### Examples

```
Ignore the permit flag added by OpenFlow.
```

```
<Sysname> system-view
[Sysname] openflow permit-flag ignore
```

## New feature: ND flood suppression

### Enabling ND flood suppression

#### About Enabling ND flood suppression

ARP or ND flood suppression reduces ARP request broadcasts or ND request multicasts by enabling the VTEP to reply to ARP or ND requests on behalf of VMs.

This feature snoops ARP or ND packets to populate the ARP or ND flood suppression table with local and remote MAC addresses. If an ARP or ND request has a matching entry, the VTEP replies to the request on behalf of the VM. If no match is found, the VTEP floods the request to both local and remote sites.

#### Procedure

1. Enter system view.  
**system-view**
2. Enter VSI view.  
**vsi** *vsi-name*
3. Enable ND flood suppression.  
**ipv6 nd suppression enable**  
By default, ND flood suppression is disabled.

### Command reference

#### display ipv6 nd suppression vsi

Use **display ipv6 nd suppression vsi** to display ND flood suppression entries.

#### Syntax

```
display ipv6 nd suppression vsi [name vsi-name] [count]
```

#### Views

Any view

#### Predefined user roles

network-admin  
network-operator

#### Parameters

**name** *vsi-name*: Specifies a VSI by its name, a case-sensitive string of 1 to 31 characters. If you do not specify a VSI, this command displays entries for all VSIs.

**slot** *slot-number*: Specifies an IRF member device by its member ID. If you do not specify a member device, this command displays entries on the master device.

**count**: Displays the number of ND flood suppression entries that match the command.

#### Examples

# Display ND flood suppression entries.

```
<Sysname> display ipv6 nd suppression vsi
```

| IPv6 address | MAC address    | VSI name | Link ID | Aging (min) |
|--------------|----------------|----------|---------|-------------|
| 1000::2      | 000f-e201-0101 | vs11     | 0x70000 | 5           |
| 1000::3      | 000f-e201-0202 | vs11     | 0x80000 | 5           |
| 1000::4      | 000f-e201-0203 | vs12     | 0x90000 | 5           |

# Display the number of ND flood suppression entries.

```
<Sysname> display ipv6 nd suppression vsi count
```

```
Total entries: 3
```

**Table 58 Command output**

| Field       | Description                                                                                                      |
|-------------|------------------------------------------------------------------------------------------------------------------|
| Link ID     | Link ID that uniquely identifies an AC or a VXLAN tunnel on a VSI.                                               |
| Aging (min) | Remaining lifetime (in minutes) of the ND flood suppression entry. When the timer expires, the entry is deleted. |

## Related commands

```
ipv6 nd suppression enable
```

```
reset ipv6 nd suppression vsi
```

## ipv6 nd suppression enable

Use **ipv6 nd suppression enable** to enable ND flood suppression.

Use **undo ipv6 nd suppression enable** to disable ND flood suppression.

## Syntax

```
ipv6 nd suppression enable
```

```
undo ipv6 nd suppression enable
```

## Default

ND flood suppression is disabled.

## Views

VSI view

## Predefined user roles

network-admin

## Usage guidelines

ND flood suppression reduces ND request multicasts by enabling the VTEP to reply to ND requests on behalf of user terminals.

This feature snoops ND packets to populate the ND flood suppression table with local and remote MAC addresses. If an ND request has a matching entry, the VTEP replies to the request on behalf of the user terminal. If no match is found, the VTEP floods the request to both local and remote sites.

## Examples

# Enable ND flood suppression for VSI **vs11**.

```
<Sysname> system-view
```

```
[Sysname] vsi vs11
```

```
[Sysname-vsi-vs11] ipv6 nd suppression enable
```

## Related commands

```
display ipv6 nd suppression vsi
reset ipv6 nd suppression vsi
```

## reset ipv6 nd suppression vsi

Use **reset ipv6 nd suppression vsi** to clear ND flood suppression entries on VSIs.

## Syntax

```
reset ipv6 nd suppression vsi [name vsi-name]
```

## Views

User view

## Predefined user roles

network-admin

## Parameters

**name vsi-name**: Specifies a VSI by its name, a case-sensitive string of 1 to 31 characters. If you do not specify a VSI, this command clears ND flood suppression entries on all VSIs.

## Examples

```
Clear ND flood suppression entries on all VSIs.
<Sysname> reset ipv6 nd suppression vsi
This command will delete all entries. Continue? [Y/N]:y
```

## Related commands

```
display ipv6 nd suppression vsi
ipv6 nd suppression enable
```

# New feature: IP address configuration of the VTEPs in a DR system

## Specifying the IP addresses of the VTEPs in a DR system

### About EVPN distributed relay

EVPN distributed relay virtualizes two VTEPs or EVPN gateways into one DR system to avoid single points of failure. The VTEPs or EVPN gateways use a virtual VTEP address to establish VXLAN tunnels to remote devices.

An AC that is attached to only one of the VTEPs in a DR system is called a single-armed AC. To ensure that the traffic of a single-armed AC is forwarded to its attached VTEP, specify the IP addresses of the VTEPs in the DR system by using the **evpn drni local** command. After you configure this command, each VTEP in a DR system changes the next hop of the routes for single-armed ACs to its local VTEP IP address when advertising the routes. When a VTEP receives BGP EVPN routes from the peer VTEP IP address specified by using this command, it does not set up a VXLAN tunnel to the peer VTEP.

You must execute the **evpn drni local** command if single-armed ACs are attached to a DR system that uses an Ethernet aggregate link as the IPL. You do not need to execute this command on a DR system that uses a VXLAN tunnel as the IPL. In such a DR system, a VTEP uses the source IP address of the IPL as the next hop of routes for single-armed ACs to ensure correct traffic forwarding.

## Procedure

1. Enter system view.  
**system-view**
2. Enable EVPN distributed relay and specify the virtual VTEP address.  
**evpn drni group** *virtual-vtep-ip*  
By default, EVPN distributed relay is disabled.  
To modify the virtual VTEP address, you must first delete the original virtual VTEP address.
3. Specify the IP addresses of the VTEPs in the DR system.  
**evpn drni local** *local-ip* **remote** *remote-ip*  
By default, the IP addresses of the VTEPs in a DR system are not specified.  
Make sure the IP address of the local VTEP belongs to a local interface. Make sure the local VTEP IP address and peer VTEP IP address are reversed on the VTEPs in the DR system.

## Command reference

### evpn drni local

Use **evpn drni local** to specify the IP addresses of the VTEPs in a DR system.

Use **undo evpn drni local** to restore the default.

#### Syntax

```
evpn drni local local-ip remote remote-ip
undo evpn drni local
```

#### Default

The IP addresses of the VTEPs in a DR system are not specified.

#### Views

System view

#### Predefined user roles

network-admin

#### Parameters

**local** *local-ip*: Specifies the IP address of the local VTEP.

**remote** *remote-ip*: Specifies the IP address of the peer VTEP.

#### Usage guidelines

An AC that is attached to only one of the VTEPs in a DR system is called a single-armed AC. After you configure this command, each VTEP in a DR system changes the next hop of the routes for single-armed ACs to its local VTEP IP address when advertising the routes. This ensures that the traffic of a single-armed AC is forwarded to its attached VTEP. When a VTEP receives BGP EVPN routes from the peer VTEP IP address specified by using this command, it does not set up a VXLAN tunnel to the peer VTEP.

You must execute this command if single-armed ACs are attached to a DR system that uses an Ethernet aggregate link as the IPL. You do not need to execute this command on a DR system that uses a VXLAN tunnel as the IPL. In such a DR system, a VTEP uses the source IP address of the IPL as the next hop of routes for single-armed ACs to ensure correct traffic forwarding.

When you execute this command, make sure the IP address of the local VTEP belongs to a local interface. Make sure the local VTEP IP address and peer VTEP IP address are reversed on the VTEPs in a DR system.

## Examples

# Specify the IP addresses of the local and peer VTEPs in the DR system as 2.2.2.2 and 3.3.3.3, respectively.

```
<Sysname> system-view
```

```
[Sysname] evpn drni local 2.2.2.2 remote 3.3.3.3
```

# New feature: Redistribution of MAC/IP advertisement routes into BGP unicast routing tables

## Redistributing MAC/IP advertisement routes into BGP unicast routing tables

### About redistributing MAC/IP advertisement routes into BGP unicast routing tables

This task enables the device to redistribute received MAC/IP advertisement routes that contain ARP or ND information into a BGP unicast routing table.

- If you perform this task for the BGP IPv4 or IPv6 unicast address family, the device will redistribute the routes into the BGP IPv4 or IPv6 unicast routing table. In addition, the device will advertise the routes to the local site.
- If you perform this task for the BGP-VPN IPv4 or IPv6 unicast address family, the device will redistribute the routes into the BGP-VPN IPv4 or IPv6 unicast routing table of the corresponding VPN instance. To advertise the routes to the local site, you must configure the **advertise 12vpn evpn** command.

### Procedure (BGP instance view)

1. Enter system view.  
**system-view**
2. Enter BGP instance view.  
**bgp as-number [ instance instance-name ]**
3. Enter BGP IPv4 or IPv6 unicast address family view.  
**address-family { ipv4 | ipv6 }**
4. Redistribute MAC/IP advertisement routes that contain ARP or ND information into the BGP IPv4 or IPv6 unicast routing table.  
**import evpn mac-ip**  
By default, MAC/IP advertisement routes that contain ARP or ND information are not redistributed into the BGP IPv4 or IPv6 unicast routing table.

### Procedure (BGP-VPN instance view)

1. Enter system view.  
**system-view**
2. Enter BGP instance view.  
**bgp as-number [ instance instance-name ]**
3. Enter BGP-VPN instance view.  
**ip vpn-instance vpn-instance-name**

4. Enter BGP-VPN IPv4 or IPv6 unicast address family view.  
**address-family { ipv4 | ipv6 }**
5. Redistribute MAC/IP advertisement routes that contain ARP or ND information into the BGP-VPN IPv4 or IPv6 unicast routing table.

**import evpn mac-ip**

By default, MAC/IP advertisement routes that contain ARP or ND information are not redistributed into the BGP-VPN IPv4 or IPv6 unicast routing table.

## Command reference

### import evpn mac-ip

Use **import evpn mac-ip** to enable the device to redistribute received MAC/IP advertisement routes that contain ARP or ND information into a BGP unicast routing table.

Use **undo import evpn mac-ip** to disable the device from redistributing received MAC/IP advertisement routes that contain ARP or ND information into a BGP unicast routing table.

#### Syntax

**import evpn mac-ip**

**undo import evpn mac-ip**

#### Default

MAC/IP advertisement routes that contain ARP or ND information are not redistributed into any BGP unicast routing table.

#### Views

BGP IPv4 unicast address family view

BGP IPv6 unicast address family view

BGP-VPN IPv4 unicast address family view

BGP-VPN IPv6 unicast address family view

#### Predefined user roles

network-admin

#### Usage guidelines

This command enables the device to redistribute received MAC/IP advertisement routes that contain ARP or ND information into a BGP unicast routing table.

- If you use this command in BGP IPv4 or IPv6 unicast address family view, the device will redistribute the routes into the BGP IPv4 or IPv6 unicast routing table. In addition, the device will advertise the routes to the local site.
- If you use this command in BGP-VPN IPv4 or IPv6 unicast address family view, the device will redistribute the routes into the BGP-VPN IPv4 or IPv6 unicast routing table of the corresponding VPN instance. To advertise the routes to the local site, you must configure the **advertise l2vpn evpn** command.

#### Examples

# Redistribute received MAC/IP advertisement routes into the BGP-VPN IPv4 unicast routing table of VPN instance **vpna**.

```
<Sysname> system-view
```

```
[Sysname] bgp 100
```

```
[Sysname-bgp-default] ip vpn-instance vpna
```

```
[Sysname-bgp-default-vpna] address-family ipv4
[Sysname-bgp-default-ipv4-vpna] import evpn mac-ip
```

## Related commands

```
advertise l2vpn evpn
```

# New feature: Route exchange between the BGP EVPN address family and the BGP VPNv4 or VPNv6 address family

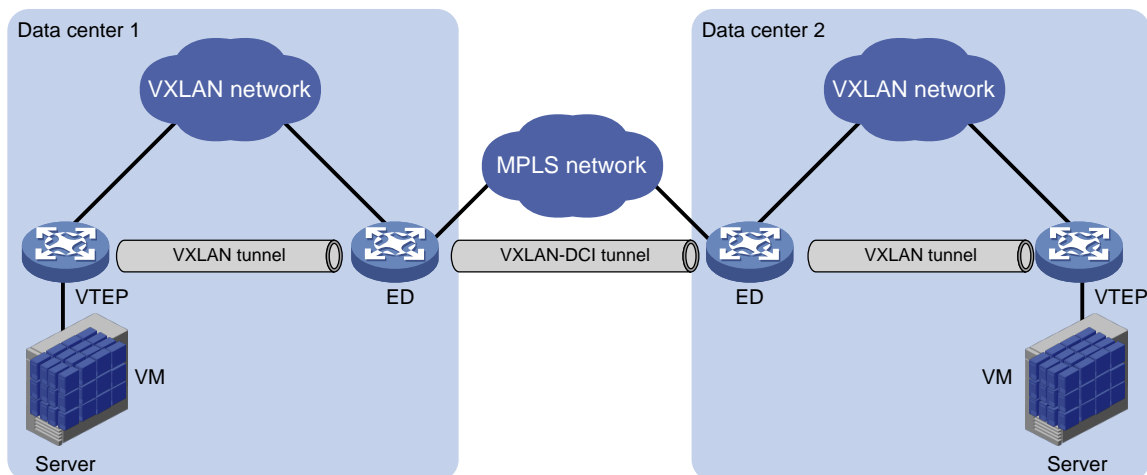
Configuring the BGP EVPN address family and the BGP VPNv4 or VPNv6 address family to exchange routes

## About route exchange

When data centers are interconnected through an MPLS L3VPN network, EVPN EDs also act as MPLS L3VPN PEs. To enable communication between the data centers, you must perform the following tasks on the EDs:

- Configure both MPLS L3VPN and EVPN.
- Configure the BGP EVPN address family and the BGP VPNv4 or VPNv6 address family to exchange routes.

**Figure 20 Data centers interconnected through an MPLS L3VPN network**



## Enabling BGP VPNv4 or VPNv6 route advertisement for the BGP EVPN address family

1. Enter system view.  
**system-view**
2. Enter BGP instance view.  
**bgp as-number [ instance instance-name ]**
3. Enter BGP EVPN address family view.  
**address-family l2vpn evpn**

4. Enable BGP VPNv4 or VPNv6 route advertisement for the BGP EVPN address family.  
**advertise l3vpn route** [ **replace-rt** ] [ **advertise-policy** *policy-name* ]  
By default, BGP VPNv4 or VPNv6 routes are not advertised through the BGP EVPN address family.  
After you execute this command, the device advertises BGP VPNv4 or VPNv6 routes as IP prefix advertisement routes through the BGP EVPN address family.

## Enabling BGP EVPN route advertisement for the BGP VPNv4 or VPNv6 address family

1. Enter system view.  
**system-view**
2. Enter BGP instance view.  
**bgp** *as-number* [ **instance** *instance-name* ]
3. Enter BGP VPNv4 address family view or BGP VPNv6 address family view.  
**address-family** { **vpn4** | **vpn6** }
4. Enable BGP EVPN route advertisement for the BGP VPNv4 or VPNv6 address family.  
**advertise evpn route** [ **replace-rt** ] [ **advertise-policy** *policy-name* ]  
By default, BGP EVPN routes are not advertised through the BGP VPNv4 or VPNv6 address family.  
After you execute this command, the device advertises IP prefix advertisement routes and MAC/IP advertisement routes that contain host route information through the BGP VPNv4 or VPNv6 address family.

## Command reference

### advertise l3vpn route

Use **advertise l3vpn route** to enable BGP VPNv4 or VPNv6 route advertisement for the BGP EVPN address family.

Use **undo advertise l3vpn route** to disable BGP VPNv4 or VPNv6 route advertisement for the BGP EVPN address family.

#### Syntax

```
advertise l3vpn route [replace-rt] [advertise-policy policy-name]
undo advertise l3vpn route
```

#### Default

BGP VPNv4 or VPNv6 routes are not advertised through the BGP EVPN address family.

#### Views

BGP EVPN address family view

#### Predefined user roles

network-admin

#### Parameters

**replace-rt**: Replaces the route targets of BGP VPNv4 or VPNv6 routes with the route targets of BGP EVPN routes. If you do not specify this keyword, route targets of BGP VPNv4 or VPNv6 routes are not modified.

**advertise-policy** *policy-name*: Specifies a routing policy to filter the BGP VPNv4 or VPNv6 routes to be advertised to BGP EVPN peers. The *policy-name* argument specifies the routing policy name, a case-sensitive string of 1 to 63 characters. If you do not specify this option, all BGP VPNv4 or VPNv6 routes are advertised to BGP EVPN peers.

## Usage guidelines

To enable communication between data centers interconnected through an MPLS L3VPN network, you must configure the BGP EVPN address family and the BGP VPNv4 or VPNv6 address family to exchange routes on EDs.

After you execute this command, the device advertises BGP VPNv4 or VPNv6 routes as IP prefix advertisement routes through the BGP EVPN address family.

## Examples

# Enable BGP VPNv4 or VPNv6 route advertisement for the BGP EVPN address family.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] address-family l2vpn evpn
[Sysname-bgp-default-evpn] advertise l3vpn route
```

## advertise evpn route

Use **advertise evpn route** to enable BGP EVPN route advertisement for the BGP VPNv4 or VPNv6 address family.

Use **undo advertise evpn route** to disable BGP EVPN route advertisement for the BGP VPNv4 or VPNv6 address family.

## Syntax

```
advertise evpn route [replace-rt] [advertise-policy policy-name]
undo advertise evpn route
```

## Default

BGP EVPN routes are not advertised through the BGP VPNv4 or VPNv6 address family.

## Views

BGP VPNv4 address family

BGP VPNv6 address family

## Predefined user roles

network-admin

## Parameters

**replace-rt**: Replaces the route targets of BGP EVPN routes with the route targets of BGP VPNv4 or VPNv6 routes. If you do not specify this keyword, route targets of BGP EVPN routes are not modified.

**advertise-policy** *policy-name*: Specifies a routing policy to filter the BGP EVPN routes to be advertised to BGP VPNv4 or VPNv6 peers. The *policy-name* argument specifies the routing policy name, a case-sensitive string of 1 to 63 characters. If you do not specify this option, all BGP EVPN routes are advertised to BGP VPNv4 or VPNv6 peers.

## Usage guidelines

To enable communication between data centers interconnected through an MPLS L3VPN network, you must configure the BGP EVPN address family and the BGP VPNv4 or VPNv6 address family to exchange routes on EDs.

After you execute this command, the device advertises IP prefix advertisement routes and MAC/IP advertisement routes that contain host route information through the BGP VPNv4 or VPNv6 address family.

## Examples

# Enable BGP EVPN route advertisement for the BGP VPNv4 address family.

```
<Sysname> system-view
[Sysname] bgp 100
[Sysname-bgp-default] address-family vpnv4
[Sysname-bgp-default-vpnv4] advertise evpn route
```

# New feature: Disabling the ACLs issued by the OVSDb controller

## Disabling the ACLs issued by the OVSDb controller

### About disabling the ACLs issued by the OVSDb controller

Perform this task on a VTEP to disable all the ACLs issued by the OVSDb controller in order to save ACL resources on the VTEP.

### Prerequisites

Before you perform this task, you must enable the OVSDb VTEP service by using the **vtep enable** command.

### Procedure

1. Enter system view.  
**system-view**
2. Disable the ACLs issued by the OVSDb controller.  
**vtep acl disable**

By default, the ACLs issued by the OVSDb controller are enabled on a VTEP.

## Command reference

### vtep acl disable

Use **vtep acl disable** to disable the ACLs issued by the OVSDb controller.

Use **undo vtep acl disable** to enable the ACLs issued by the OVSDb controller.

### Syntax

```
vtep acl disable
undo vtep acl disable
```

### Default

The ACLs issued by the OVSDb controller are enabled on the device.

### Views

System view

## Predefined user roles

network-admin

## Usage guidelines

Before you use this command, you must use the **vtep enable** command to enable the OVSDB VTEP service.

Use the **vtep acl disable** command on a VTEP to disable all the ACLs issued by the OVSDB controller in order to save ACL resources on the VTEP.

## Examples

# Disable the ACLs issued by the OVSDB controller.

```
<Sysname> system-view
```

```
[Sysname] vtep enable
```

```
[sysname] vtep acl disable
```

# New feature: Replacement of the L3 VXLAN ID and RD of IP prefix advertisement routes on EDs

## Enabling an ED to replace the L3 VXLAN ID and RD of IP prefix advertisement routes

### About replacement of the L3 VXLAN ID and RD of IP prefix advertisement routes

In an EVPN-DCI network, use this feature to enable communication between data centers that use different L3 VXLAN IDs or hide the L3 VXLAN ID of a data center. After you enable this feature on an ED, the ED performs the following operations after receiving IP prefix advertisement routes:

1. Matches the route targets of the routes with the import route targets of local VPN instances.
2. Replaces the L3 VXLAN ID and RD of the routes with the L3 VXLAN ID and RD of the matching local VPN instance.
3. Advertises the routes to a VTEP or remote ED.

### Restrictions and guidelines

After you configure this feature, an ED advertises only IP prefix advertisement routes with the replaced L3 VXLAN ID and RD. The IP prefix advertisement routes with the original L3 VXLAN ID and RD are not advertised.

If the RD of a received IP prefix advertisement route is identical to the RD of the matching local VPN instance, an ED does not replace the L3 VXLAN ID of the route or regenerate the route. As a result, the ED does not advertise the route. As a best practice, assign unique RDs to VPN instances on different EVPN gateways and EDs when you use this feature.

## Procedure

1. Enter system view.  
**system-view**
2. Enter BGP instance view.  
**bgp as-number [ instance instance-name ]**
3. Enter BGP EVPN address family view.  
**address-family l2vpn evpn**
4. Replace the L3 VXLAN ID and RD of IP prefix advertisement routes.  
**peer { group-name | ipv4-address [ mask-length ] } re-originated**

By default, the device does not modify the L3 VXLAN ID or RD of the IP prefix advertisement routes that are received from peers or peer groups.

## Command reference

### peer re-originated

Use **peer re-originated** to replace the L3 VXLAN ID and RD of IP prefix advertisement routes.

Use **undo peer re-originated** to restore the default.

#### Syntax

```
peer { group-name | ipv4-address [mask-length] } re-originated
undo peer { group-name | ipv4-address [mask-length] } re-originated
```

#### Default

The device does not modify the L3 VXLAN ID or RD of the IP prefix advertisement routes that are received from peers or peer groups.

#### Views

BGP EVPN address family view

#### Predefined user roles

network-admin

#### Parameters

*group-name*: Specifies a peer group by its name, a case-sensitive string of 1 to 47 characters. The peer group must exist.

*ipv4-address*: Specifies a peer by its IPv4 address. The peer must exist.

*mask-length*: Specifies a mask length in the range of 0 to 32. To specify a subnet, you must specify both the *ipv4-address* and *mask-length* arguments.

#### Usage guidelines

In an EVPN-DCI network, use this command to enable communication between data centers that use different L3 VXLAN IDs or hide the L3 VXLAN ID of a data center. After you execute this command on an ED, the ED performs the following operations after receiving IP prefix advertisement routes:

1. Matches the route targets of the routes with the import route targets of local VPN instances.
2. Replaces the L3 VXLAN ID and RD of the routes with the L3 VXLAN ID and RD of the matching local VPN instance.
3. Advertises the routes to a VTEP or remote ED.

After you execute this command, an ED advertises only IP prefix advertisement routes with the replaced L3 VXLAN ID and RD. The IP prefix advertisement routes with the original L3 VXLAN ID and RD are not advertised.

If the RD of a received IP prefix advertisement route is identical to the RD of the matching local VPN instance, an ED does not replace the L3 VXLAN ID of the route or regenerate the route. As a result, the ED does not advertise the route. As a best practice, assign unique RDs to VPN instances on different EVPN gateways and EDs when you use this command.

#### Examples

# Replace the L3 VXLAN ID and RD of IP prefix advertisement routes.

```
<Sysname> system-view
[Sysname] bgp 100
```

```
[Sysname-bgp-default] address-family l2vpn evpn
[Sysname-bgp-default-evpn] peer 1.1.1.1 re-originated
```

## New feature: BGP EVPN route advertisement suppression

### Suppressing BGP EVPN route advertisement

#### About BGP EVPN route advertisement suppression

To reduce the number of BGP EVPN routes on EDs of an EVPN-DCI network, suppress the advertisement of specific BGP EVPN routes on the EDs.

#### Restrictions and guidelines

If two VSI interfaces on EVPN gateways of different data centers use the same IP address, do not suppress the advertisement of MAC/IP advertisement routes on the EDs of the data centers. If you suppress the advertisement of these routes, the EDs cannot communicate with each other.

#### Procedure

1. Enter system view.  
**system-view**
  2. Enter BGP instance view.  
**bgp** *as-number* [ **instance** *instance-name* ]
  3. Enter BGP EVPN address family view.  
**address-family** **l2vpn evpn**
  4. Suppress the advertisement of specific BGP EVPN routes to a peer or peer group.  
**peer** { *group-name* | *ipv4-address* [ *mask-length* ] } **advertise evpn-route suppress** { **ip-prefix** | **mac-ip** }
- By default, advertisement of BGP EVPN routes is not suppressed.

### Command reference

#### peer advertise evpn-route suppress

Use **peer advertise evpn-route suppress** to suppress the advertisement of specific BGP EVPN routes to a peer or peer group.

Use **undo peer advertise evpn-route suppress** to restore the default.

#### Syntax

```
peer { group-name | ipv4-address [mask-length] } advertise evpn-route suppress { ip-prefix | mac-ip }
undo peer { group-name | ipv4-address [mask-length] } advertise evpn-route suppress { ip-prefix | mac-ip }
```

#### Default

Advertisement of BGP EVPN routes is not suppressed.

#### Views

BGP EVPN address family view

## Predefined user roles

network-admin

## Parameters

*group-name*: Specifies a peer group by its name, a case-sensitive string of 1 to 47 characters. The peer group must exist.

*ipv4-address*: Specifies a peer by its IPv4 address. The peer must exist.

*mask-length*: Specifies a mask length in the range of 0 to 32. To specify a subnet, you must specify both the *ipv4-address* and *mask-length* arguments.

**ip-prefix**: Suppresses IP prefix advertisement routes.

**mac-ip**: Suppresses MAC/IP advertisement routes.

## Usage guidelines

To reduce the number of BGP EVPN routes on EDs of an EVPN-DCI network, suppress the advertisement of specific BGP EVPN routes on the EDs.

If two VSI interfaces on EVPN gateways of different data centers use the same IP address, do not suppress the advertisement of MAC/IP advertisement routes on the EDs of the data centers. If you suppress the advertisement of these routes, the EDs cannot communicate with each other.

## Examples

# Suppress the IP prefix advertisement routes advertised to peer 1.1.1.1.

```
<Sysname> system-view
```

```
[Sysname] bgp 100
```

```
[Sysname-bgp-default] address-family l2vpn evpn
```

```
[Sysname-bgp-default-evpn] peer 1.1.1.1 advertise evpn-route suppress ip-prefix
```

# New feature: Removal of the default-gateway extended community attribute from EVPN gateway routes

## Removing the default-gateway extended community attribute from EVPN gateway routes

About removal of the default-gateway extended community attribute from EVPN gateway routes

An EVPN gateway route is the route for the IP address of a VSI interface on an EVPN gateway. The EVPN gateway routes advertised by an EVPN gateway contain the default-gateway extended community attribute. The EVPN gateway routes with that attribute cannot be used as ECMP routes. You can perform this task to remove the default-gateway extended community attribute from EVPN gateway routes for the routes to be used for load sharing.

## Procedure

1. Enter system view.  
**system-view**
2. Enter BGP instance view.  
**bgp** *as-number* [ **instance** *instance-name* ]
3. Enter BGP EVPN address family view.  
**address-family** *l2vpn evpn*

4. Remove the default-gateway extended community attribute from the EVPN gateway routes advertised to a peer or peer group.

```
peer { group-name | ipv4-address [mask-length] } default-gateway
no-advertise
```

By default, EVPN gateway routes advertised to peers and peer groups contain the default-gateway extended community attribute.

## Command reference

### peer default-gateway no-advertise

Use **peer default-gateway no-advertise** to remove the default-gateway extended community attribute from the EVPN gateway routes advertised to a peer or peer group.

Use **undo peer default-gateway no-advertise** to restore the default.

#### Syntax

```
peer { group-name | ipv4-address [mask-length] } default-gateway
no-advertise

undo peer { group-name | ipv4-address [mask-length] } default-gateway
no-advertise
```

#### Default

EVPN gateway routes advertised to peers and peer groups contain the default-gateway extended community attribute.

#### Views

BGP EVPN address family view

#### Predefined user roles

network-admin

#### Parameters

*group-name*: Specifies a peer group by its name, a case-sensitive string of 1 to 47 characters. The peer group must exist.

*ipv4-address*: Specifies a peer by its IPv4 address. The peer must exist.

*mask-length*: Specifies a mask length in the range of 0 to 32. To specify a subnet, you must specify both the *ipv4-address* and *mask-length* arguments.

#### Usage guidelines

An EVPN gateway route is the route for the IP address of a VSI interface on an EVPN gateway. By default, the EVPN gateway routes advertised by an EVPN gateway contain the default-gateway extended community attribute. The EVPN gateway routes with that attribute cannot be used as ECMP routes. You can use this command to remove the default-gateway extended community attribute from EVPN gateway routes for the routes to be used for load sharing.

#### Examples

```
Remove the default-gateway extended community attribute from the EVPN gateway routes
advertised to peer 1.1.1.1.
```

```
<Sysname> system-view
```

```
[Sysname] bgp 100
```

```
[Sysname-bgp-default] address-family l2vpn evpn
```

```
[Sysname-bgp-default-evpn] peer 1.1.1.1 default-gateway no-advertise
```

# New feature: Configuring a 16/32-Gbps FC interface

## Configuring a 16/32-Gbps FC interface

### About configuring a 16/32-Gbps FC interface

Before this software version, the HPE FlexFabric 5940 2-slot Switch (JH397A) and HPE FlexFabric 5940 4-slot Switch (JH398A) support only 8-Gbps FC interfaces after they are installed with the LSWM124XG2QFC (JH184A) expansion module. From this software version, the HPE FlexFabric 5940 2-slot Switch (JH397A) and HPE FlexFabric 5940 4-slot Switch (JH398A) support the LSWM116FC (JH339A) expansion module. By default, all interfaces on the expansion module are 32-Gbps FC interfaces. The interfaces can be split into 16-Gbps FC interfaces or converted to 10-GE interfaces.

This section describes only configurations specific to 16/32-Gbps FC interfaces. Configurations supported by 8-Gbps FC interfaces are also supported by 16/32-Gbps FC interfaces.

### Configuring the speed of a 16/32-Gbps FC interface

#### Restrictions and guidelines

As a best practice for a 16/32-Gbps FC interface to work correctly, do not configure the speed as autonegotiation.

#### Procedure

1. Enter system view.  
**system-view**
  2. Enter 16/32-Gbps FC interface view.  
**interface fc** *interface-number*
  3. Set the speed of the FC interface.  
**speed** { 8000 | 16000 | 32000 | auto }
- By default, the speed of an FC interface is autonegotiated.

### Splitting a 32-Gbps FC interface and combining 16-Gbps FC breakout interfaces

#### About 32-Gbps FC interface splitting and 16-Gbps FC breakout interface combining

By default, all interfaces on the LSWM116FC expansion module are 32-Gbps FC interfaces. A 32-Gbps FC interface is composed of two FC interfaces, for example, FC 1/1/1 and FC 1/1/2, FC 1/1/3 and FC 1/1/4. The FC interface ending with an even number can work correctly. The FC interface ending with an odd number cannot be identified by the device.

You can use a 32-Gbps FC interface as a single interface or split a 32-Gbps FC interface into two 16-Gbps FC breakout interfaces. For example, 32-Gbps FC interface FC 1/1/2 can be split into two 16-Gbps breakout interfaces FC 1/1/1 and FC 1/1/2. In this case, the 16-Gbps breakout interface ending with the odd number can also be identified by the device and forward packets. You can use the **using thirty-twogigfc** command in the view of either FC 1/1/1 or FC 1/1/2 to combine the two 16-Gbps breakout interfaces into 32-Gbps FC interface FC 1/1/2.

By splitting a 32-Gbps FC interface into two 16-Gbps FC breakout interfaces, you can improve port density, reduce costs, and improve networking flexibility. The 16-Gbps FC breakout interfaces support the same configuration and attributes as common 16-Gbps interfaces.

By combining two 16-Gbps FC breakout interfaces into a 32-Gbps FC interface, you can get higher bandwidth on a single interface. The 32-Gbps FC interface supports the same configuration and attributes as common 32-Gbps interfaces.

### Splitting a 32-Gbps FC interface into two 16-Gbps FC breakout interfaces

1. Enter system view.  
**system-view**
2. Enter 32-Gbps FC interface view.  
**interface fc** *interface-number*
3. Split the 32-Gbps FC interface into two 16-Gbps FC breakout interfaces.  
**using sixteengigfc**  
By default, an FC interface is a 32-Gbps interface.  
This command deletes the 32-Gbps FC interface and creates two 16-Gbps FC breakout interfaces.

### Combining two 16-Gbps FC breakout interfaces into a 32-Gbps FC interface

1. Enter system view.  
**system-view**
2. Enter the view of either 16-Gbps FC breakout interface.  
**interface fc** *interface-number*
3. Combine the two 16-Gbps FC breakout interfaces into a 32-Gbps FC interface.  
**using thirty-twogigfc**  
By default, an FC interface is a 32-Gbps interface.  
This command deletes the two 16-Gbps FC breakout interfaces and creates a 32-Gbps FC interface.  
The 32-Gbps FC interface has the same number as the 16-Gbps FC breakout interface that ends with an even number.

## Changing the port type between a Layer 2 Ethernet interface and an FC interface

### About changing the port type between a Layer 2 Ethernet interface and an FC interface

A 32-Gbps FC interface can be changed to two 10-GE Layer 2 Ethernet interfaces in the view of the 32-Gbps FC interface. For example, you can change interface FC 1/1/2 to interfaces Ten-GigabitEthernet 1/1/1 and Ten-GigabitEthernet 1/1/2 in the view of FC 1/1/2.

You can also change the two 16-Gbps breakout interface of a 32-Gbps FC interface to two 10-GE Layer 2 Ethernet interfaces in the view of either breakout interface. For example, you can change breakout interfaces FC 1/1/1 and FC 1/1/2 of 32-Gbps FC interface FC 1/1/2 to interfaces Ten-GigabitEthernet 1/1/1 and Ten-GigabitEthernet 1/1/2 in the view of breakout interface FC 1/1/1 or FC 1/1/2.

You can change the two 10-GE Layer 2 Ethernet interfaces back to a 32-Gbps FC interface in the view of either 10-GE Layer 2 Ethernet interface.

### Procedure

For the procedure or commands used to changing the port type, see the FC interface chapter in the FC and FCoE configuration guide or command reference of the switch.

# Command reference

## New command: using

Use **using** to split a 32-Gbps FC interface into two 16-Gbps FC breakout interfaces or combine two 16-Gbps FC breakout interfaces back into a 32-Gbps FC interface.

### Syntax

```
using { sixteengigfc | thirty-twogigfc }
```

### Default

An FC interface is a 32-Gbps interface.

### Views

FC interface view

### Predefined user roles

network-admin

### Parameters

**sixteengigfc**: Splits a 32-Gbps FC interface into two 16-Gbps FC breakout interfaces.

**thirty-twogigfc**: Combines two 16-Gbps FC breakout interfaces split from a 32-Gbps FC interface into a 32-Gbps FC interface.

### Usage guidelines

To improve port density, reduce costs, and improve network flexibility, you can split a 32-Gbps FC interface into two 16-Gbps breakout interfaces. The 16-Gbps breakout interfaces support the same configuration and attributes as common 16-Gbps FC interfaces.

To get higher bandwidth on a single interface, you can combine two 16-Gbps FC breakout interfaces back into a 32-Gbps FC interface in the view of either 16-Gbps FC breakout interface. The 32-Gbps FC interface supports the same configuration and attributes as common 32-Gbps FC interfaces.

You can use the **display this** command in FC interface view to identify 32-Gbps FC interfaces and 16-Gbps FC breakout interfaces. If the **using thirty-twogigfc** attribute appears, the interface is a 32-Gbps FC interface. If the **using sixteengigfc** attribute appears, the interface is a 16-Gbps FC breakout interface.

### Examples

# Split 32-Gbps FC interface FC 1/1/2 into two 16-Gbps FC breakout interfaces.

```
<System> system-view
[System] interface fc 1/1/2
[System-Fc1/1/2] using sixteengigfc
The interface 32G FC1/1/2 will be deleted. Continue? [Y/N]:y
```

# Combine two 16-Gbps FC breakout interfaces FC 1/1/1 and FC 1/1/2 back into a 32-Gbps FC interface.

```
<System> system-view
[System] interface fc 1/1/1
[System-Fc1/1/1] using thirty-twogigfc
The interfaces 16G FC1/1/1 and 16G FC1/1/2 will be deleted. Continue? [Y/N]:y
```

## Modified command: speed

Use **speed** to set the speed of an FC interface.

Use **undo speed** to restore the default.

### Old syntax

**speed** { 2000 | 4000 | 8000 | auto }

### New syntax

**speed** { 2000 | 4000 | 8000 | 16000 | 32000 | auto }

### Default

The speed of an FC interface is autonegotiated.

### Views

FC interface view

### Change description

From this software version, the switch supports 16/32-Gbps FC interfaces. The **16000** and **32000** parameters were added to the **speed** command. When configuring this command, follow these restrictions and guidelines:

- 32-Gbps FC interfaces support speeds of only 8000 Mbps, 16000 Mbps, 32000 Mbps, and autonegotiation.
- 16-Gbps FC interfaces support speeds of only 8000 Mbps, 16000 Mbps, and autonegotiation.
- 8-Gbps FC interfaces support speeds of only 2000 Mbps, 4000 Mbps, 8000 Mbps, and autonegotiation.
- As a best practice for a 16/32-Gbps FC interface to work correctly, do not configure the speed as autonegotiation.

## New feature: Enabling area ID-to-F\_Port binding

### Enabling area ID-to-F\_Port binding

#### About area ID-to-F\_Port binding

This feature affects the area ID assignment for N\_Ports on all F\_Ports and the mask length of related direct routes.

After this feature is enabled, all N\_Ports are logged out and log in again. The N\_Ports that log in through the same F\_Port are assigned the same area ID. The switch generates one direct route with a 16-bit mask for an F\_Port that is an FC interface or FC aggregate interface. If the F\_Port is a VFC interface, the switch generates one direct route with a 24-bit mask for each N\_Port on the F\_Port.

#### Restrictions and guidelines

Before enabling this feature, correctly plan FC addresses for N\_Ports.

#### Procedure

1. Enter system view.  
**system-view**
2. Enable area ID-to-F\_Port binding.  
**fc port-bind-area enable**

By default, area ID-to-F\_Port binding is disabled.

### fc port-bind-area enable

Use **fc port-bind-area enable** to enable area ID-to-F\_Port binding.

Use **undo fc port-bind-area enable** to disable area ID-to-F\_Port binding.

### Syntax

```
fc port-bind-area enable
undo fc port-bind-area enable
```

### Default

Area ID-to-F\_Port binding is disabled.

### Views

System view

### Predefined user roles

network-admin

### Usage guidelines

Only FCF and FCF-NPV switches support this command.

This command affects the area ID assignment for N\_Ports on all F\_Ports and the mask length of related direct routes.

After this command is configured, the switch logs out all N\_Ports that have logged in to an F\_Port and deletes the persistent FC address entries for these N\_Ports. The switch assigns an FC address to the first N\_Port that logs in to the F\_Port in the way used when area ID-to-F\_Port binding is not enabled. Subsequent N\_Ports that log in to the same F\_Port use the area ID of the first N\_Port. For an F\_Port that is an FC interface or FC aggregate interface, the switch generates only one direct route with a 16-bit mask. For an F\_Port that is a VFC interface, the switch generates one direct route with a 24-bit mask for each N\_Port.

If the area ID of the bound or desired FC address for an N\_Port is different from the bound area ID for the F\_Port, the N\_Port cannot log in. Before enabling this command, correctly plan FC addresses for N\_Ports.

### Examples

```
Enable area ID-to-F_Port binding.
```

```
<Sysname> system-view
```

```
[Sysname] fc port-bind-area enable
```

```
All nodes will be logged out, and FC ID persistence data will be deleted. Continue? [Y/N]:y
```

### Related commands

```
fc port-bind-area enable
```

### display fc port-bind-area

Use **display fc port-bind-area** to display information about the area ID-to-F\_Port binding feature.

### Syntax

```
display fc port-bind-area [vsan vsan-id]
```

### Views

Any view

### Predefined user roles

network-admin

network-operator

## Parameters

**vsan** *vsan-id*: Specifies a VSAN by its ID in the range of 1 to 3839. If you do not specify a VSAN, this command displays information about the area ID-to-F\_Port binding feature for all VSANs. On an FCF-NPV switch, this command displays information about the area ID-to-F\_Port binding feature for only VSANs operating in FCF mode.

## Usage guidelines

Only FCF and FCF-NPV switches support this command.

## Examples

# Display information about the area ID-to-F\_Port binding feature for VSAN 1.

```
<Sysname> display fc port-bind-area vsan 10
```

```
FC port-bind-area function: Enabled
```

```
VSAN 10:
```

```
Interface Area ID
```

```
Fc1/0/1 0xa0
```

```
Vfc1 0xa1
```

# Display information about the area ID-to-F\_Port binding feature for VSAN 1.

```
<Sysname> display fc port-bind-area vsan 10
```

```
FC port-bind-area function: Disabled
```

**Table 59 Command output**

| Field                      | Description                                                    |
|----------------------------|----------------------------------------------------------------|
| FC port-bind-area function | Area ID-to-F_Port binding feature status: Enabled or Disabled. |
| Interface                  | F_Port that has logged-in N_Ports.                             |
| Area ID                    | Area ID bound to the F_Port.                                   |

## Related commands

```
fc port-bind-area enable
```

# New feature: BB\_Credit recovery for FC interfaces

## Enabling BB\_Credit recovery for an FC interface

### About BB\_Credit recovery

This feature allows two FC interfaces to recover the connection when BB\_Credits are lost.

### Restrictions and guidelines

For this feature to take effect, you must enable it on both FC interfaces.

### Procedure

1. Enter system view.  
**system-view**
2. Enter FC interface view.  
**interface fc** *interface-number*
3. Enable BB\_Credit recovery for the FC interface.

**fc2bcredit recovery enable**

By default, BB\_Credit recovery is disabled for an FC interface.

## Command reference

### fc2bcredit recovery enable

Use **fc2bcredit recovery enable** to enable BB\_Credit recovery for an FC interface.

Use **undo fc2bcredit recovery enable** to disable BB\_Credit recovery for an FC interface.

#### Syntax

**fc2bcredit recovery enable**

**undo fc2bcredit recovery enable**

#### Default

BB\_Credit recovery is disabled for an FC interface.

#### Views

FC interface view

#### Predefined user roles

network-admin

#### Usage guidelines

This command allows two FC interfaces to recover the connection when BB\_Credits are lost.

For the feature to take effect, you must configure the command on both FC interfaces.

#### Examples

```
Enable BB_Credit recovery for FC 1/0/1.
<Sysname> system-view
[Sysname] interface fc 1/0/1
[Sysname-Fc1/0/1] fc2bcredit recovery enable
```

## Modified feature: Software upgrade

### Feature change description

Due to memory limitations, switches labeled with the following product codes cannot ensure correct operation of F2701 and later versions:

- HPE FlexFabric 5940 48SFP+ 6QSFP+ Switch (JH395A)
- HPE FlexFabric 5940 48XGT 6QSFP+ Switch (JH394A)

If you load F2701 or a later version on these switches, the following error message and log message are displayed:

- Error message: The memory space of the IRF member device in slot x is insufficient for the software version to run correctly. Please upgrade the switch to Release 26xx or earlier.
- Log message: DRV/3/HARDWARE\_INCOMPATIBLE: The memory space of the IRF member device in slot x is insufficient for the software version to run correctly.

## Command changes

None.

## Modified feature: Using issu commands for an ISSU

### Feature change description

You can specify patch image files for the **display version comp-matrix**, **issu load**, and **issu one-step** commands.

## Command changes

### Modified command: display version comp-matrix

#### Old syntax

```
display version comp-matrix file { boot filename | system filename |
feature filename<1-30> } *

display version comp-matrix file ipe ipe-filename
```

#### New syntax

```
display version comp-matrix file { boot filename | system filename |
feature filename<1-30> | patch filename<1-16> } *

display version comp-matrix file ipe ipe-filename [patch
filename<1-16>]
```

#### Views

Any view

#### Parameters

**patch**: Specifies a space-separated list of up to 16 patch image files.

**filename**: Specifies a .bin file in the *filesystemname/filename.bin* format. The file must be stored in the root directory of a file system on the device. The value string excluding the file system location section (if any) can have a maximum of 63 characters. For more information about specifying a file, see file system management in *Fundamentals Configuration Guide*.

#### Change description

Before modification: The command does not support specifying patch image files.

After modification: The command supports specifying patch image files.

### Modified command: issu load

#### Old syntax

```
issu load file { boot filename | system filename | feature filename<1-30> }
* slot slot-number<1-9> [reboot]

issu load file ipe ipe-filename slot slot-number<1-9> [reboot]
```

## New syntax

```
issu load file { boot filename | system filename | feature filename&<1-30>
| patch filename&<1-16> } * slot slot-number&<1-9> [reboot]

issu load file ipe ipe-filename [patch filename&<1-16>] slot
slot-number&<1-9> [reboot]
```

## Views

User view

## Parameters

**patch**: Specifies a space-separated list of up to 16 patch image files.

*filename*: Specifies a .bin file in the *filesystemname/filename.bin* format. The file must be stored in the root directory of a file system on the device. The value string excluding the file system location section (if any) can have a maximum of 63 characters. For more information about specifying a file, see file system management in *Fundamentals Configuration Guide*.

## Change description

Before modification: The command does not support specifying patch image files.

After modification: The command supports specifying patch image files.

## Modified command: issu one-step

### Old syntax

```
issu one-step file { boot filename | system filename | feature
filename&<1-30> } * [slot slot-number&<1-9>] [reboot]

issu one-step file ipe ipe-filename slot slot-number&<1-9> [reboot]
```

### New syntax

```
issu one-step file { boot filename | system filename | feature
filename&<1-30> | patch filename&<1-16> } * [slot slot-number&<1-9>]
[reboot]

issu one-step file ipe ipe-filename [patch filename&<1-16>] slot
slot-number&<1-9> [reboot]
```

## Views

User view

## Parameters

**patch**: Specifies a space-separated list of up to 16 patch image files.

*filename*: Specifies a .bin file in the *filesystemname/filename.bin* format. The file must be stored in the root directory of a file system on the device. The value string excluding the file system location section (if any) can have a maximum of 63 characters. For more information about specifying a file, see file system management in *Fundamentals Configuration Guide*.

## Change description

Before modification: The command does not support specifying patch image files.

After modification: The command supports specifying patch image files.

## Modified feature: Per-packet load sharing mode of aggregation groups

### Feature change description

Per-packet load sharing mode was removed for aggregate interfaces.

### Command changes

Modified command: link-aggregation load-sharing mode

#### Old syntax

```
link-aggregation load-sharing mode { { destination-ip | destination-mac |
source-ip | source-mac } * | flexible | per-packet }
```

#### New syntax

```
link-aggregation load-sharing mode { { destination-ip | destination-mac |
source-ip | source-mac } * | flexible }
```

#### Views

Layer 2 aggregate interface view

Layer 3 aggregate interface view

#### Change description

The **per-packet** keyword was removed.

## Modified feature: Configuring MAC-to-VLAN entries

### Feature change description

From this release, the 802.1p priority parameter in a MAC-to-VLAN entry was changed into dot1p.

### Command changes

Modified command: mac-vlan mac-address

#### Old syntax

```
mac-vlan mac-address mac-address [mask mac-mask] vlan vlan-id [dot1q priority]
```

#### New syntax

```
mac-vlan mac-address mac-address [mask mac-mask] vlan vlan-id [dot1p priority]
```

#### Views

System view

#### Change description

Before modification: The 802.1p priority parameter is dot1q.

After modification: The 802.1p priority parameter is dot1p.

# Modified feature: Configuring the advertisable TLVs

## Feature change description

This release added support for advertising the link aggregation TLV in the 802.3 organizationally specific TLV set on an interface.

## Command changes

### Modified command: lldp tlv-enable

#### Old syntax

In Layer 2 Ethernet interface view:

- For nearest bridge agents:

```
lldp tlv-enable { basic-tlv { all | port-description |
system-capability | system-description | system-name |
management-address-tlv [ipv6] [ip-address | interface loopback
interface-number] } | dot1-tlv { all | congestion-notification |
port-vlan-id | link-aggregation | dcbx | protocol-vlan-id [vlan-id] |
vlan-name [vlan-id] | management-vid [mvlan-id] } | dot3-tlv { all |
mac-physic | max-frame-size | power } | med-tlv { all | capability |
inventory | network-policy [vlan-id] | power-over-ethernet |
location-id { civic-address device-type country-code { ca-type
ca-value } &<1-10> | elin-address tel-number } } }

undo lldp tlv-enable { basic-tlv { all | port-description |
system-capability | system-description | system-name |
management-address-tlv [ipv6] [ip-address | interface loopback
interface-number] } | dot1-tlv { all | congestion-notification |
port-vlan-id | link-aggregation | dcbx | protocol-vlan-id | vlan-name |
management-vid } | dot3-tlv { all | mac-physic | max-frame-size | power }
| med-tlv { all | capability | inventory | network-policy [vlan-id] |
power-over-ethernet | location-id } }
```
- For nearest non-TPMR bridge agents:

```
lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ipv6] [ip-address] } |
dot1-tlv { all | congestion-notification | evb | port-vlan-id |
link-aggregation } }

undo lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ipv6] [ip-address] } |
dot1-tlv { all | congestion-notification | evb | port-vlan-id |
link-aggregation } }
```
- For nearest customer bridge agents:

```
lldp agent nearest-customer tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ipv6] [ip-address] } |
dot1-tlv { all | congestion-notification | port-vlan-id |
link-aggregation } }

undo lldp agent nearest-customer tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
```

```

system-name | management-address-tlv [ipv6] [ip-address] } |
dot1-tlv { all | congestion-notification | port-vlan-id |
link-aggregation } }

```

In Layer 3 Ethernet interface view:

```

lldp tlv-enable { basic-tlv { all | port-description | system-capability |
system-description | system-name | management-address-tlv [ipv6]
[ip-address | interface loopback interface-number] } | dot1-tlv { all |
link-aggregation } | dot3-tlv { all | mac-physic | max-frame-size | power } |
med-tlv { all | capability | inventory | power-over-ethernet | location-id
{ civic-address device-type country-code { ca-type ca-value } &<1-10> |
elin-address tel-number } } }

```

```

lldp agent { nearest-nontpmr | nearest-customer } tlv-enable { basic-tlv
{ all | port-description | system-capability | system-description |
system-name | management-address-tlv [ipv6] [ip-address] } | dot1-tlv
{ all | link-aggregation } }

```

```

undo lldp tlv-enable { basic-tlv { all | port-description |
system-capability | system-description | system-name |
management-address-tlv [ipv6] [ip-address | interface loopback
interface-number] } | dot1-tlv { all | link-aggregation } | dot3-tlv { all |
mac-physic | max-frame-size | power } | med-tlv { all | capability |
inventory | power-over-ethernet | location-id } }

```

```

undo lldp agent { nearest-nontpmr | nearest-customer } tlv-enable
{ basic-tlv { all | port-description | system-capability |
system-description | system-name | management-address-tlv [ipv6]
[ip-address] } | dot1-tlv { all | link-aggregation } }

```

In management Ethernet interface view:

```

lldp tlv-enable { basic-tlv { all | port-description | system-capability |
system-description | system-name | management-address-tlv [ipv6]
[ip-address] } | dot1-tlv { all | link-aggregation } | dot3-tlv { all |
mac-physic | max-frame-size | power } | med-tlv { all | capability |
inventory | power-over-ethernet | location-id { civic-address device-type
country-code { ca-type ca-value } &<1-10> | elin-address tel-number } } }

```

```

lldp agent { nearest-nontpmr | nearest-customer } tlv-enable { basic-tlv
{ all | port-description | system-capability | system-description |
system-name | management-address-tlv [ipv6] [ip-address] } | dot1-tlv
{ all | link-aggregation } }

```

```

undo lldp tlv-enable { basic-tlv { all | port-description |
system-capability | system-description | system-name |
management-address-tlv [ipv6] [ip-address] } | dot1-tlv { all |
link-aggregation } | dot3-tlv { all | mac-physic | max-frame-size | power } |
med-tlv { all | capability | inventory | power-over-ethernet |
location-id } }

```

```

undo lldp agent { nearest-nontpmr | nearest-customer } tlv-enable
{ basic-tlv { all | port-description | system-capability |
system-description | system-name | management-address-tlv [ipv6]
[ip-address] } | dot1-tlv { all | link-aggregation } }

```

## New syntax

In Layer 2 Ethernet interface view:

- For nearest bridge agents:

```
lldp tlv-enable { basic-tlv { all | port-description |
system-capability | system-description | system-name |
management-address-tlv [ipv6] [ip-address | interface loopback
interface-number] } | dot1-tlv { all | congestion-notification |
port-vlan-id | link-aggregation | dcbx | protocol-vlan-id [vlan-id] |
vlan-name [vlan-id] | management-vid [mvlan-id] } | dot3-tlv { all |
link-aggregation | mac-physic | max-frame-size | power } | med-tlv { all
| capability | inventory | network-policy [vlan-id] |
power-over-ethernet | location-id { civic-address device-type
country-code { ca-type ca-value } <1-10> | elin-address
tel-number } } }
```

```
undo lldp tlv-enable { basic-tlv { all | port-description |
system-capability | system-description | system-name |
management-address-tlv [ipv6] [ip-address | interface loopback
interface-number] } | dot1-tlv { all | congestion-notification |
port-vlan-id | link-aggregation | dcbx | protocol-vlan-id | vlan-name |
management-vid } | dot3-tlv { all | link-aggregation | mac-physic |
max-frame-size | power } | med-tlv { all | capability | inventory |
network-policy [vlan-id] | power-over-ethernet | location-id }
```

- For nearest non-TPMR bridge agents:

```
lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ipv6] [ip-address] } |
dot1-tlv { all | congestion-notification | evb | port-vlan-id |
link-aggregation } | dot3-tlv { all | link-aggregation } }
```

```
undo lldp agent nearest-nontpmr tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ipv6] [ip-address] } |
dot1-tlv { all | congestion-notification | evb | port-vlan-id |
link-aggregation } | dot3-tlv { all | link-aggregation } }
```

- For nearest customer bridge agents:

```
lldp agent nearest-customer tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ipv6] [ip-address] } |
dot1-tlv { all | congestion-notification | port-vlan-id |
link-aggregation } | dot3-tlv { all | link-aggregation } }
```

```
undo lldp agent nearest-customer tlv-enable { basic-tlv { all |
port-description | system-capability | system-description |
system-name | management-address-tlv [ipv6] [ip-address] } |
dot1-tlv { all | congestion-notification | port-vlan-id |
link-aggregation } | dot3-tlv { all | link-aggregation } }
```

In Layer 3 Ethernet interface view:

```
lldp tlv-enable { basic-tlv { all | port-description | system-capability |
system-description | system-name | management-address-tlv [ipv6]
[ip-address | interface loopback interface-number] } | dot1-tlv { all |
link-aggregation } | dot3-tlv { all | link-aggregation | mac-physic |
max-frame-size | power } | med-tlv { all | capability | inventory |
power-over-ethernet | location-id { civic-address device-type
country-code { ca-type ca-value } <1-10> | elin-address tel-number } } }
```

```
lldp agent { nearest-nontpmr | nearest-customer } tlv-enable { basic-tlv
{ all | port-description | system-capability | system-description |
system-name | management-address-tlv [ipv6] [ip-address] } | dot1-tlv
{ all | link-aggregation } | dot3-tlv { all | link-aggregation } }
```

```
undo lldp tlv-enable { basic-tlv { all | port-description |
system-capability | system-description | system-name |
management-address-tlv [ipv6] [ip-address | interface loopback
interface-number] } | dot1-tlv { all | link-aggregation } | dot3-tlv { all |
link-aggregation | mac-physic | max-frame-size | power } | med-tlv { all |
capability | inventory | power-over-ethernet | location-id } }

undo lldp agent { nearest-nontpmr | nearest-customer } tlv-enable
{ basic-tlv { all | port-description | system-capability |
system-description | system-name | management-address-tlv [ipv6]
[ip-address] } | dot1-tlv { all | link-aggregation } | dot3-tlv { all |
link-aggregation } }
```

In management Ethernet interface view:

```
lldp tlv-enable { basic-tlv { all | port-description | system-capability |
system-description | system-name | management-address-tlv [ipv6]
[ip-address] } | dot1-tlv { all | link-aggregation } | dot3-tlv { all |
link-aggregation | mac-physic | max-frame-size | power } | med-tlv { all |
capability | inventory | power-over-ethernet | location-id { civic-address
device-type country-code { ca-type ca-value }<1-10> | elin-address
tel-number } } }
```

```
lldp agent { nearest-nontpmr | nearest-customer } tlv-enable { basic-tlv
{ all | port-description | system-capability | system-description |
system-name | management-address-tlv [ipv6] [ip-address] } | dot1-tlv
{ all | link-aggregation } | dot3-tlv { all | link-aggregation } }
```

```
undo lldp tlv-enable { basic-tlv { all | port-description |
system-capability | system-description | system-name |
management-address-tlv [ipv6] [ip-address] } | dot1-tlv { all |
link-aggregation } | dot3-tlv { all | link-aggregation | mac-physic |
max-frame-size | power } | med-tlv { all | capability | inventory |
power-over-ethernet | location-id } }
```

```
undo lldp agent { nearest-nontpmr | nearest-customer } tlv-enable
{ basic-tlv { all | port-description | system-capability |
system-description | system-name | management-address-tlv [ipv6]
[ip-address] } | dot1-tlv { all | link-aggregation } | dot3-tlv { all |
link-aggregation } }
```

## Change description

Before modification: Advertisement of the link aggregation TLV in the 802.3 organizationally specific TLV set is not supported on an interface.

After modification: Advertisement of the link aggregation TLV in the 802.3 organizationally specific TLV set is supported on an interface.

# Modified feature: Specifying the management address advertised in global management address TLV advertisement setting

## Feature change description

This release added support for specifying the IP address of an M-GigabitEthernet interface as the management address advertised in global management address TLV advertisement setting.

## Command changes

Modified command: `lldp global tlv-enable basic-tlv management-address-tlv`

### Old syntax

```
lldp [agent { nearest-customer | nearest-nontpmr }] global tlv-enable
basic-tlv management-address-tlv [ipv6] { ip-address | interface loopback
interface-number | interface vlan-interface interface-number }
```

### New syntax

```
lldp [agent { nearest-customer | nearest-nontpmr }] global tlv-enable
basic-tlv management-address-tlv [ipv6] { ip-address | interface loopback
interface-number | interface m-gigabitethernet interface-number |
interface vlan-interface interface-number }
```

### Views

System view

### Change description

The `interface m-gigabitethernet interface-number` option was added to the command. You can specify the IP address of an M-GigabitEthernet interface as the management address to be advertised in global management address TLV advertisement setting.

## Modified feature: Configuring aging timer for dynamic ARP entries

### Feature change description

From this release, support for setting the aging timer of dynamic ARP entries was added to the following views:

- Layer 3 Ethernet interface view.
- Layer 3 Ethernet subinterface view.
- Layer 3 aggregate interface view.
- Layer 3 aggregate subinterface view.
- VLAN interface view.
- VSI interface view.

## Command changes

Modified command: `arp timer aging`

### Syntax

```
arp timer aging { aging-minutes | second aging-seconds }
undo arp timer aging
```

### Views

System view

Layer 3 Ethernet interface view

Layer 3 Ethernet subinterface view  
Layer 3 aggregate interface view  
Layer 3 aggregate subinterface view  
VSI interface view  
VLAN interface view

### Change description

Before modification: Setting the aging timer of dynamic ARP entries is only supported in system view.

After modification: Setting the aging timer of dynamic ARP entries is supported in the following views:

- System view.
- Layer 3 Ethernet interface view.
- Layer 3 Ethernet subinterface view.
- Layer 3 aggregate interface view.
- Layer 3 aggregate subinterface view.
- VSI interface view.
- VLAN interface view.

## Modified feature: Enabling ARP snooping

### Feature change description

From this release, support for enabling ARP snooping was added to VSI view.

### Command changes

#### Modified command: arp snooping enable

##### Syntax

```
arp snooping enable
undo arp snooping enable
```

##### Views

VLAN view  
VSI view

### Change description

Before modification: The ARP snooping can only be enabled in VLAN view.

After modification: The ARP snooping can be enabled in VLAN view and VSI view.

# Modified feature: Displaying ARP snooping entries

## Feature change description

From this release, the following changes were made to the **display arp snooping** command:

- Support displaying ARP snooping entries in a VSI.
- The **vlan** keyword is required in the command for displaying ARP snooping entries in a VLAN.  
Before modification: **display arp snooping** [ **vlan** *vlan-id* ] [ **slot** *slot-number* ] [ **count** ].  
After modification: **display arp snooping vlan** [ *vlan-id* ] [ **slot** *slot-number* ] [ **count** ].
- The **vlan** keyword is added to the **display arp snooping vlan ip** command.  
Before modification: **display arp snooping ip** *ip-address* [ **slot** *slot-number* ].  
After modification: **display arp snooping vlan ip** *ip-address* [ **slot** *slot-number* ].

## Command changes

### Modified command: display arp snooping

#### Old syntax

```
display arp snooping [vlan vlan-id] [slot slot-number] [count]
display arp snooping ip ip-address [slot slot-number]
```

#### New syntax

```
display arp snooping { vlan [vlan-id] | vsi [vsi-name] } [slot slot-number]
[count]
display arp snooping vlan ip ip-address [slot slot-number]
```

#### Views

Any view

#### Predefined user roles

network-admin

#### Change description

Before modification:

- The *vlan-id* argument is required when you specify the **vlan** keyword for the **display arp snooping** command.
- The **vlan** keyword is not in the **display arp snooping ip** command.
- Displaying ARP snooping entries in a VSI is not supported.

After modification:

- The *vlan-id* argument is optional when you specify the **vlan** keyword for the **display arp snooping** command.
- The **vlan** keyword is added to the **display arp snooping vlan ip** command.
- Support displaying ARP snooping entries in a VSI.

# Modified feature: Clearing ARP snooping entries

## Feature change description

From this release, the following changes were made to the **reset arp snooping** command:

- Support clearing ARP snooping entries in a VSI.
- The **reset arp snooping [ ip ip-address | vlan vlan-id ]** command was changed to the following commands, where the **vlan** keyword became a must in each command:
  - **reset arp snooping vlan [ vlan-id ]**.
  - **reset arp snooping vlan ip ip-address**.

## Command changes

### Modified command: reset arp snooping

#### Old syntax

```
reset arp snooping [ip ip-address | vlan vlan-id]
```

#### New syntax

```
reset arp snooping { vlan [vlan-id] | vsi [vsi-name] }
reset arp snooping vlan ip ip-address
```

#### Views

User view

#### Predefined user roles

network-admin

#### Change description

Before modification:

- Clearing ARP snooping entries in a VSI is not supported.
- The **vlan** keyword is optional for the **reset arp snooping [ ip ip-address | vlan vlan-id ]** command.

After modification:

- Support clearing ARP snooping entries in a VSI.
- The **vlan** keyword is required for the **reset arp snooping vlan [ vlan-id ]** and **reset arp snooping vlan ip ip-address** commands.

# Modified feature: Configuring DHCP snooping handling strategy for Option 82 in request messages

## Feature change description

The DHCP snooping device supports setting the handling strategy to **append** for processing Option 82 in request messages.

## Command changes

### Modified command: dhcp snooping information strategy

#### Old syntax

```
dhcp snooping information strategy { drop | keep | replace }
```

#### New syntax

```
dhcp snooping information strategy { append | drop | keep | replace }
```

#### Parameters

**append**: Processes a DHCP message as follows:

- If the DHCP message does not carry Option 82, the device forwards the message after adding the Option 82 according to the padding configuration.
- If the DHCP message carries Option 82, the device processes the message as follows:
  - Forwards the message after padding the Vendor-Specific sub-option with the content specified in the **dhcp snooping information vendor-specific** command.
  - Forwards the message without changing Option 82 if the **dhcp snooping information vendor-specific** command is not configured.

#### Change description

Before modification: The DHCP snooping device supports the following handling strategies for Option 82 in request messages: **drop**, **keep**, and **replace**.

After modification: The DHCP snooping device supports the following handling strategies for Option 82 in request messages: **append**, **drop**, **keep**, and **replace**.

## Modified feature: Setting the interface MTU for IPv6 packets

### Feature change description

This release added support of setting the interface MTU for IPv6 packets in VSI interface view.

## Command changes

### Modified command: ipv6 mtu

#### Syntax

```
ipv6 mtu size
```

```
undo ipv6 mtu
```

#### Change description

Before modification: Setting the interface MTU for IPv6 packets is not supported in VSI interface view.

After modification: Setting the interface MTU for IPv6 packets is supported in VSI interface view.

# Modified feature: Support for specifying next hop indexes in static routes

## Feature change description

From this release, you can specify a next hop index when configuring a static route.

## Command changes

### Modified command: ip route-static

#### Old syntax

```
ip route-static { dest-address { mask-length | mask } | group group-name } { interface-type interface-number [next-hop-address] [backup-interface interface-type interface-number [backup-nexthop backup-nexthop-address] [permanent] | bfd { control-packet | echo-packet } | permanent | track track-entry-number] | next-hop-address [recursive-lookup host-route] [bfd control-packet bfd-source ip-address | permanent | track track-entry-number] | vpn-instance d-vpn-instance-name next-hop-address [recursive-lookup host-route] [bfd control-packet bfd-source ip-address | permanent | track track-entry-number] } [preference preference] [tag tag-value] [description text]
```

```
ip route-static vpn-instance s-vpn-instance-name { dest-address { mask-length | mask } | group group-name } { interface-type interface-number [next-hop-address] [backup-interface interface-type interface-number [backup-nexthop backup-nexthop-address] [permanent] | bfd { control-packet | echo-packet } | permanent | track track-entry-number] | next-hop-address [recursive-lookup host-route] [public] [bfd control-packet bfd-source ip-address | permanent | track track-entry-number] | vpn-instance d-vpn-instance-name next-hop-address [recursive-lookup host-route] [bfd control-packet bfd-source ip-address | permanent | track track-entry-number] } [preference preference] [tag tag-value] [description text]
```

#### New syntax

```
ip route-static { dest-address { mask-length | mask } | group group-name } { interface-type interface-number [next-hop-address [nexthop-index index-string]] [backup-interface interface-type interface-number [backup-nexthop backup-nexthop-address]] [permanent] | bfd { control-packet | echo-packet } | permanent | track track-entry-number] | next-hop-address [nexthop-index index-string] [recursive-lookup host-route] [bfd control-packet bfd-source ip-address | permanent | track track-entry-number] | vpn-instance d-vpn-instance-name next-hop-address [recursive-lookup host-route] [bfd control-packet bfd-source ip-address | permanent | track track-entry-number] } [preference preference] [tag tag-value] [description text]
```

```
ip route-static vpn-instance s-vpn-instance-name { dest-address { mask-length | mask } | group group-name } { interface-type interface-number [next-hop-address [nexthop-index index-string]] [backup-interface interface-type interface-number [backup-nexthop backup-nexthop-address]] [permanent] | bfd { control-packet | echo-packet } | permanent | track track-entry-number] | next-hop-address [nexthop-index index-string] [recursive-lookup host-route] [public] [bfd control-packet bfd-source ip-address | permanent | track track-entry-number] | vpn-instance d-vpn-instance-name next-hop-address [recursive-lookup host-route] [bfd control-packet bfd-source ip-address | permanent | track track-entry-number] } [preference preference] [tag tag-value] [description text]
```

## Views

System view

## Parameters

**nexthop-index** *index-string*: Specifies the index of the next hop. The *index-string* argument represents the index, which is a case-insensitive string of 1 to 93 characters. If you do not specify an index for the next hop, the static route uses the output interface and the IP address of the next hop as the next hop index (for example, gigabitethernet1/0/12.2.2.2). Do not specify the same next hop index for routes to the same destination.

## Change description

Before modification: The command does not support specifying the index of the next hop.

After modification: The command supports using the **nexthop-index** *index-string* option to specify the index of the next hop.

# Modified feature: Support for configuring OSPF commands in VSI interface view

## Feature change description

From this release, you can configure the following OSPF commands in VSI interface view.

- `ospf area`
- `ospf authentication-mode`
- `ospf bfd enable`
- `ospf cost`
- `ospf database-filter`
- `ospf dr-priority`
- `ospf fast-reroute lfa-backup`
- `ospf mtu-enable`
- `ospf network-type`
- `ospf packet-size`
- `ospf prefix-suppression`
- `ospf primary-path-detect bfd`
- `ospf timer dead`
- `ospf timer hello`
- `ospf timer poll`
- `ospf timer retransmit`
- `ospf trans-delay`
- `ospf ttl-security`

## Modified feature: Displaying IS-IS LSP log information

### Feature change description

The **refreshed** and **purged** keywords were added to the **display isis event-log lsp** command for you to view log information about purged LSPs or refreshed LSPs.

### Command changes

Modified command: **display isis event-log lsp**

#### Old syntax

```
display isis event-log lsp [level-1 | level-2] [process-id]
```

#### New syntax

```
display isis event-log lsp { purged | refreshed } [level-1 | level-2]
[process-id]
```

#### Change description

Before modification: The **refreshed** and **purged** keywords are not supported.

After modification: The **refreshed** and **purged** keywords are supported for you to view log information about purged LSPs or refreshed LSPs.

## Modified feature: Clearing IS-IS LSP log information

### Feature change description

The **refreshed** and **purged** keywords were added to the **reset isis event-log lsp** command for you to clear log information about purged LSPs or refreshed LSPs.

### Command changes

Modified command: **reset isis event-log lsp**

#### Old syntax

```
reset isis event-log lsp [process-id]
```

#### New syntax

```
reset isis event-log lsp { purged | refreshed } [process-id]
```

#### Change description

Before modification: The **refreshed** and **purged** keywords are not supported.

After modification: The **refreshed** and **purged** keywords are supported for you to clear log information about purged LSPs or refreshed LSPs.

# Modified feature: Filtering advertised/received BGP routes

## Feature change description

In this release, you can specify an ACL by its name to filter advertised or received BGP routes.

## Command changes

### Modified command: filter-policy export/filter-policy import

#### Old syntax

In BGP IPv4 unicast address family view/BGP-VPN IPv4 unicast address family view/BGP VPNv4 address family view/BGP IPv4 multicast address family view:

```
filter-policy { ipv4-acl-number | prefix-list ipv4-prefix-list-name }
export [direct | { isis | ospf | rip } process-id | static]
```

In BGP IPv6 unicast address family view/BGP-VPN IPv6 unicast address family view/BGP VPNv6 address family view/BGP IPv6 multicast address family view:

```
filter-policy { ipv6-acl-number | prefix-list ipv6-prefix-list-name }
export [direct | { isisv6 | ospfv3 | ripng } process-id | static]
```

In BGP IPv4 unicast address family view/BGP-VPN IPv4 unicast address family view/BGP VPNv4 address family view/BGP IPv4 multicast address family view:

```
filter-policy { ipv4-acl-number | prefix-list ipv4-prefix-list-name }
import
```

In BGP IPv6 unicast address family view/BGP-VPN IPv6 unicast address family view/BGP VPNv6 address family view/BGP IPv6 multicast address family view:

```
filter-policy { ipv6-acl-number | prefix-list ipv6-prefix-list-name }
import
```

#### New syntax

In BGP IPv4 unicast address family view/BGP-VPN IPv4 unicast address family view/BGP VPNv4 address family view/BGP IPv4 multicast address family view:

```
filter-policy { ipv4-acl-number | name ipv4-acl-name | prefix-list
ipv4-prefix-list-name } export [direct | { isis | ospf | rip } process-id |
static]
```

In BGP IPv6 unicast address family view/BGP-VPN IPv6 unicast address family view/BGP VPNv6 address family view/BGP IPv6 multicast address family view:

```
filter-policy { ipv6-acl-number | name ipv6-acl-name | prefix-list
ipv6-prefix-list-name } export [direct | { isisv6 | ospfv3 | ripng }
process-id | static]
```

In BGP IPv4 unicast address family view/BGP-VPN IPv4 unicast address family view/BGP VPNv4 address family view/BGP IPv4 multicast address family view:

```
filter-policy { ipv4-acl-number | name ipv4-acl-name | prefix-list
ipv4-prefix-list-name } import
```

In BGP IPv6 unicast address family view/BGP-VPN IPv6 unicast address family view/BGP VPNv6 address family view/BGP IPv6 multicast address family view:

```
filter-policy { ipv6-acl-number | name ipv6-acl-name | prefix-list
ipv6-prefix-list-name } import
```

## Views

BGP IPv4 unicast address family view  
 BGP-VPN IPv4 unicast address family view  
 BGP VPNv4 address family view  
 BGP IPv6 unicast address family view  
 BGP-VPN IPv6 unicast address family view  
 BGP VPNv6 address family view  
 BGP IPv4 multicast address family view  
 BGP IPv6 multicast address family view

## Change description

Before modification: The **name** *ipv4-acl-name* and **name** *ipv6-acl-name* options are not supported.

After modification: The **name** *ipv4-acl-name* and **name** *ipv6-acl-name* options are supported for you to specify an ACL by its name.

# Modified feature: Filtering routes advertised to or received from a peer or peer group by using an ACL

## Feature change description

In this release, you can specify an ACL by its name to filter routes advertised to or received from a peer or peer group.

## Command changes

### Modified command: peer filter-policy

#### Old syntax

In BGP IPv4 unicast address family view/BGP-VPN IPv4 unicast address family view/BGP VPNv4 address family view/BGP IPv4 multicast address family view:

```
peer { group-name | ipv4-address [mask-length] } filter-policy
ipv4-acl-number{ export | import }
```

```
undo peer { group-name | ipv4-address [mask-length] } filter-policy
[ipv4-acl-number]{ export | import }
```

In BGP IPv6 unicast address family view:

```
peer { group-name | ipv4-address [mask-length] | ipv6-address
[prefix-length]} filter-policy ipv6-acl6-number{ export | import }
```

```
undo peer { group-name | ipv4-address [mask-length] | ipv6-address
[prefix-length]} filter-policy[ipv6-acl-number]{ export | import }
```

In BGP-VPN IPv6 unicast address family view/BGP IPv6 multicast address family view:

```
peer { group-name | ipv6-address [prefix-length] } filter-policy
ipv6-acl-number{export | import}
```

```
undo peer { group-name | ipv6-address [prefix-length] } filter-policy
[ipv6-acl-number]{export | import}
```

In BGP VPNv6 address family view:

```
peer { group-name | ipv4-address [mask-length] } filter-policy
ipv6-acl-number{export | import}
```

```
undo peer { group-name | ipv4-address [mask-length] } filter-policy
[ipv6-acl-number]{export | import}
```

## New syntax

In BGP IPv4 unicast address family view/BGP-VPN IPv4 unicast address family view/BGP VPNv4 address family view/BGP IPv4 multicast address family view:

```
peer { group-name | ipv4-address [mask-length] } filter-policy
{ ipv4-acl-number | name ipv4-acl-name } { export | import }
```

```
undo peer { group-name | ipv4-address [mask-length] } filter-policy
[ipv4-acl-number | name ipv4-acl-name] { export | import }
```

In BGP IPv6 unicast address family view:

```
peer { group-name | ipv4-address [mask-length] | ipv6-address
[prefix-length] } filter-policy { ipv6-acl-number | name ipv6-acl-name }
{ export | import }
```

```
undo peer { group-name | ipv4-address [mask-length] | ipv6-address
[prefix-length] } filter-policy [ipv6-acl-number | name ipv6-acl-name]
{ export | import }
```

In BGP-VPN IPv6 unicast address family view/BGP IPv6 multicast address family view:

```
peer { group-name | ipv6-address [prefix-length] } filter-policy
{ ipv6-acl-number | name ipv6-acl-name } { export | import }
```

```
undo peer { group-name | ipv6-address [prefix-length] } filter-policy
[ipv6-acl-number | name ipv6-acl-name] { export | import }
```

In BGP VPNv6 address family view:

```
peer { group-name | ipv4-address [mask-length] } filter-policy
{ ipv6-acl-number | name ipv6-acl-name } { export | import }
```

```
undo peer { group-name | ipv4-address [mask-length] } filter-policy
[ipv6-acl-number | name ipv6-acl-name] { export | import }
```

## Views

BGP IPv4 unicast address family view

BGP-VPN IPv4 unicast address family view

BGP VPNv4 address family view

BGP IPv6 unicast address family view

BGP-VPN IPv6 unicast address family view

BGP VPNv6 address family view

BGP IPv4 multicast address family view

BGP IPv6 multicast address family view

## Change description

Before modification: The **name** *ipv4-acl-name* and **name** *ipv6-acl-name* options are not supported.

After modification: The **name** *ipv4-acl-name* and **name** *ipv6-acl-name* options are supported for you to specify an ACL by its name.

# Modified feature: Configuring BMP

## Feature change description

Support for the following features was added to BGP BMP:

- Specifying the source interface of TCP connections to the BMP server.
- Sending routes advertised to the monitored peer or peer group to the BMP server.
- Sending the optimal routes in the routing table to the BMP server.

## Command changes

### New command: server connect-interface

Use **server connect-interface** to specify the source interface of TCP connections to the BMP server.

Use **undo server connect-interface** to restore the default.

### Syntax

```
server connect-interface interface-type interface-number
undo server connect-interface
```

### Default

BGP uses the primary IP address of the output interface in the optimal route to the BMP server as the source address of TCP connections to the BMP server.

### Views

BMP server view

### Predefined user roles

network-admin

### Parameters

*interface-type interface-number*: Specifies an interface by its type and number.

### Usage guidelines

After you configure this command, BGP uses the address of the specified interface to establish TCP connections to the BMP server.

Do not specify a virtual template (VT) interface for this command.

### Examples

# Configure Loopback 0 as the source interface of TCP connections to BMP server 5.

```
<Sysname> system-view
```

```
[Sysname] bmp server 5
```

```
[Sysname-bmpserver-5] server address 100.1.1.1 port 8888
```

```
[Sysname-bmpserver-5] server connect-interface loopback0
```

## New command: route-mode adj-rib-out

Use **route-mode adj-rib-out** to configure BGP to send routes advertised to the monitored peer or peer group to the BMP server.

Use **undo route-mode adj-rib-out** to restore the default.

### Syntax

```
route-mode adj-rib-out
undo route-mode adj-rib-out
```

### Default

BGP does not send routes advertised to the monitored peer or peer group to the BMP server.

### Views

BMP server view

### Predefined user roles

network-admin

### Usage guidelines

To configure BGP to send routes advertised to the monitored peer or peer group to the BMP server, make sure BGP has established a TCP connection to the BMP server.

### Examples

# Configure BGP to send routes advertised to the monitored peer or peer group to BMP server 5.

```
<Sysname> system-view
[Sysname] bmp server 5
[Sysname-bmpserver-5] route-mode adj-rib-out
```

## New command: route-mode loc-rib

Use **route-mode loc-rib** to configure BGP to send the optimal routes in the routing table to the BMP server.

Use **undo route-mode loc-rib** to restore the default.

### Syntax

```
route-mode loc-rib
undo route-mode loc-rib
```

### Default

BGP does not send the optimal routes in the routing table to the BMP server.

### Views

BMP server view

### Predefined user roles

network-admin

### Usage guidelines

To configure BGP to send the optimal routes in the routing table to the BMP server, make sure BGP has established a TCP connection to the BMP server.

This command takes effect only on BGP IPv4 unicast and BGP-VPN IPv4 unicast routes.

## Examples

# Configure BGP to send the optimal routes in the routing table to BMP server 5.

```
<Sysname> system-view
[Sysname] bmp server 5
[Sysname-bmpserver-5] route-mode loc-rib
```

# Modified feature: Displaying detailed BGP routing information

## Feature change description

From this release, the **delay** route state can be displayed in the output from the following commands:

- `display bgp routing-table ipv4 multicast`
- `display bgp routing-table ipv4 rtfilter`
- `display bgp routing-table ipv4 unicast`
- `display bgp routing-table ipv6 multicast`
- `display bgp routing-table ipv6 unicast`

## Command changes

### Modified command: display bgp routing-table ipv4 multicast

#### Syntax

```
display bgp [instance instance-name] routing-table ipv4 multicast
[ipv4-address [{ mask-length | mask } [longest-match]] | ipv4-address
[mask-length | mask] advertise-info | as-path-acl as-path-acl-number |
community-list { { basic-community-list-number | comm-list-name }
[whole-match] | adv-community-list-number } | peer ipv4-address
{ advertised-routes | received-routes } [ipv4-address [mask-length | mask]
| statistics] | statistics]
```

#### Views

Any view

#### Change description

After modification, the **delay** route state can be displayed in the command output, as shown in the following example:

# Display detailed information about BGP IPv4 multicast routes destined to network 5.5.5.5/32.

```
<Sysname> display bgp routing-table ipv4 multicast 5.5.5.5 32
```

```
BGP local router ID: 192.168.1.139
```

```
Local AS number: 100
```

```
Paths: 1 available, 1 best
```

```
BGP routing table information of 5.5.5.5/32:
```

```

From : 192.168.1.62 (192.168.1.62)
Rely nexthop : 192.168.1.62
Original nexthop: 192.168.1.62
OutLabel : NULL
RxPathID : 0x0
TxPathID : 0x0
AS-path : (null)
Origin : incomplete
Attribute value : MED 0, localpref 100, pref-val 0
State : valid, internal, best, delay
IP precedence : N/A
QoS local ID : N/A
Traffic index : N/A

```

**Table 60 Command output**

| Field            | Description                                                                                                                                                                                                                                                                                                                              |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Paths            | Number of routes: <ul style="list-style-type: none"> <li><b>available</b>—Number of valid routes.</li> <li><b>best</b>—Number of optimal routes.</li> </ul>                                                                                                                                                                              |
| From             | IP address of BGP peer that advertised the route.                                                                                                                                                                                                                                                                                        |
| Rely Nexthop     | Next hop found by route recursion. If no next hop is found, this field displays <b>not resolved</b> .                                                                                                                                                                                                                                    |
| Original nexthop | Original next hop of the route. If the route was obtained from a BGP update message, the original next hop is the next hop IP address in the message.                                                                                                                                                                                    |
| OutLabel         | Outgoing label of the route.                                                                                                                                                                                                                                                                                                             |
| RxPathID         | Add-Path ID of received routes.                                                                                                                                                                                                                                                                                                          |
| TxPathID         | Add-Path ID of advertised routes.                                                                                                                                                                                                                                                                                                        |
| AS-path          | AS_PATH attribute of the route, which records the ASs the route has passed and avoids routing loops.                                                                                                                                                                                                                                     |
| Origin           | Origin of the route: <ul style="list-style-type: none"> <li><b>igp</b>—Originated in the AS. The origin of routes advertised with the <b>network</b> command is IGP.</li> <li><b>egp</b>—Learned through EGP.</li> <li><b>incomplete</b>—Unknown origin. The origin of routes redistributed from IGP protocols is INCOMPLETE.</li> </ul> |
| Attribute value  | BGP path attributes: <ul style="list-style-type: none"> <li><b>MED</b>—MED value.</li> <li><b>localpref</b>—Local preference value.</li> <li><b>pref-val</b>—Preferred value.</li> <li><b>pre</b>—Route preference.</li> </ul>                                                                                                           |

| Field         | Description                                                                                                                                                                                                                                                                                                                                          |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| State         | Current state of the route: <ul style="list-style-type: none"> <li>• valid.</li> <li>• internal.</li> <li>• external.</li> <li>• local.</li> <li>• synchronize.</li> <li>• best.</li> <li>• delay. (The value means optimal route selection will be delayed for this route. This field is displayed only in detailed output information.)</li> </ul> |
| IP precedence | IP precedence in the range of 0 to 7. <b>N/A</b> indicates that the route does not support this field.                                                                                                                                                                                                                                               |
| QoS local ID  | QoS local ID in the range of 1 to 4095. <b>N/A</b> indicates that the route does not support this field.                                                                                                                                                                                                                                             |
| Traffic index | Traffic index in the range of 1 to 64. <b>N/A</b> indicates that the route does not support this field.                                                                                                                                                                                                                                              |

Modified command: display bgp routing-table ipv4 rtfilter

### Syntax

```
display bgp [instance instance-name] routing-table ipv4 rtfilter
[default-rt [advertise-info] | [origin-as as-number] [route-target
[advertise-info]] | peer ipv4-address { advertised-routes |
received-routes } [default-rt | [origin-as as-number] [route-target] |
statistics] | statistics]
```

### Views

Any view

### Change description

After modification, the **delay** route state can be displayed in the command output, as shown in the following example:

# Display detailed information about BGP IPv4 RT filter route 100:1.

```
<Sysname> display bgp routing-table ipv4 rtfilter 100:1
```

```
BGP local router ID: 192.168.100.1
```

```
Local AS number: 100
```

```
Origin AS: 100
```

```
Total number of routes: 1
```

```
Paths: 1 available, 1 best
```

```
BGP routing table information of <100:1>/96:
```

```
Imported route.
```

```
Original nexthop: 10.2.1.1
```

```
OutLabel : NULL
```

```
RxPathID : 0x0
```

```
TxPathID : 0x0
```

```
AS-path : (null)
```

```

Origin : igp
Attribute value : MED 0, pref-val 32768, pre 0
State : valid, local, best, delay
IP precedence : N/A
QoS local ID : N/A
Traffic index : N/A

```

**Table 61 Command output**

| Field            | Description                                                                                                                                                                                                                                                                                                                              |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Paths            | Number of routes: <ul style="list-style-type: none"> <li><b>available</b>—Number of valid routes.</li> <li><b>best</b>—Number of optimal routes.</li> </ul>                                                                                                                                                                              |
| Imported route   | The BGP RT filter route is locally generated.                                                                                                                                                                                                                                                                                            |
| Original nexthop | Original next hop of the route. If the route was obtained from a BGP update message, the original next hop is the next hop IP address in the message.                                                                                                                                                                                    |
| OutLabel         | Outgoing label of the route.                                                                                                                                                                                                                                                                                                             |
| RxPathID         | Add-Path ID of received routes.                                                                                                                                                                                                                                                                                                          |
| TxPathID         | Add-Path ID of advertised routes.                                                                                                                                                                                                                                                                                                        |
| AS-path          | AS_PATH attribute of the route, which records the ASs the route has passed and avoids routing loops.                                                                                                                                                                                                                                     |
| Origin           | Origin of the route: <ul style="list-style-type: none"> <li><b>igp</b>—Originated in the AS. The origin of routes advertised with the <b>network</b> command is IGP.</li> <li><b>egp</b>—Learned through EGP.</li> <li><b>incomplete</b>—Unknown origin. The origin of routes redistributed from IGP protocols is INCOMPLETE.</li> </ul> |
| Attribute value  | BGP path attributes: <ul style="list-style-type: none"> <li><b>MED</b>—MED value.</li> <li><b>localpref</b>—Local preference value.</li> <li><b>pref-val</b>—Preferred value.</li> <li><b>pre</b>—Route preference.</li> </ul>                                                                                                           |
| State            | Current state of the route: <ul style="list-style-type: none"> <li>valid.</li> <li>internal.</li> <li>external.</li> <li>local.</li> <li>synchronize.</li> <li>best.</li> <li>delay. (The value means optimal route selection will be delayed for this route. This field is displayed only in detailed output information.)</li> </ul>   |
| From             | IP address of BGP peer that advertised the route.                                                                                                                                                                                                                                                                                        |
| Rely Nexthop     | Next hop found by route recursion. If no next hop is found, this field displays <b>not resolved</b> .                                                                                                                                                                                                                                    |
| IP precedence    | IP precedence in the range of 0 to 7. <b>N/A</b> indicates that the route does not support this field.                                                                                                                                                                                                                                   |
| QoS local ID     | QoS local ID in the range of 1 to 4095. <b>N/A</b> indicates that the route                                                                                                                                                                                                                                                              |

|               |                                                                                                         |
|---------------|---------------------------------------------------------------------------------------------------------|
|               | does not support this field.                                                                            |
| Traffic index | Traffic index in the range of 1 to 64. <b>N/A</b> indicates that the route does not support this field. |
| Backup route  | The route is a backup route.                                                                            |

## Modified command: display bgp routing-table ipv4 unicast

### Syntax

```
display bgp [instance instance-name] routing-table ipv4 [unicast]
[vpn-instance vpn-instance-name] [ipv4-address [{ mask-length | mask }
[longest-match]] | ipv4-address [mask-length | mask] advertise-info |
as-path-acl as-path-acl-number | community-list
{ { basic-community-list-number | comm-list-name } [whole-match] |
adv-community-list-number } | peer ipv4-address { advertised-routes |
received-routes } [ipv4-address [mask-length | mask] | statistics] |
statistics]
```

### Views

Any view

### Change description

After modification, the **delay** route state can be displayed in the command output, as shown in the following examples:

# Display detailed information about BGP IPv4 unicast routes destined to network 10.2.1.0/24.

```
<Sysname> display bgp routing-table ipv4 10.2.1.0 24
```

```
BGP local router ID: 192.168.100.1
```

```
Local AS number: 100
```

```
Paths: 2 available, 1 best
```

```
BGP routing table information of 10.2.1.0/24:
```

```
Imported route.
```

```
Original nexthop: 10.2.1.1
```

```
OutLabel : NULL
```

```
RxPathID : 0x0
```

```
TxPathID : 0x0
```

```
AS-path : (null)
```

```
Origin : igp
```

```
Attribute value : MED 0, pref-val 0, pre 0
```

```
State : valid, local, best, delay
```

```
IP precedence : N/A
```

```
QoS local ID : N/A
```

```
Traffic index : N/A
```

```
From : 10.2.1.2 (192.168.100.2)
```

```
Rely nexthop : not resolved
```

```
Original nexthop: 10.2.1.2
```

```
OutLabel : NULL
```

RxPathID : 0x0  
TxPathID : 0x0  
AS-path : 200  
Origin : igp  
Attribute value : MED 0, pref-val 0, pre 255  
State : external  
IP precedence : N/A  
QoS local ID : N/A  
Traffic index : N/A

# Display detailed information about the BGP IPv4 unicast route destined to address 1.1.1.1/32.

<Sysname> display bgp routing-table ipv4 1.1.1.1 32

BGP local router ID: 192.168.100.1  
Local AS number: 100

Paths: 2 available, 1 best

BGP routing table information of 1.1.1.1/32:

From : 10.2.1.1 (192.168.100.3)  
Rely nexthop : 10.2.1.1  
Original nexthop: 10.2.1.1  
OutLabel : NULL  
RxPathID : 0x0  
TxPathID : 0x0  
AS-path : (null)  
Origin : igp  
Attribute value : MED 0, pref-val 0, pre 0  
State : valid, local, best, delay  
IP precedence : N/A  
QoS local ID : N/A  
Traffic index : N/A

Backup route.

From : 10.2.1.2 (192.168.100.2)  
Rely nexthop : 10.2.1.2  
Original nexthop: 10.2.1.2  
OutLabel : NULL  
RxPathID : 0x0  
TxPathID : 0x0  
AS-path : 200  
Origin : igp  
Attribute value : MED 0, pref-val 0, pre 255  
State : external  
IP precedence : N/A  
QoS local ID : N/A  
Traffic index : N/A

**Table 62 Command output**

| Field            | Description                                                                                                                                                                                                                                                                                                                                          |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Paths            | Number of routes: <ul style="list-style-type: none"> <li>• <b>available</b>—Number of valid routes.</li> <li>• <b>best</b>—Number of optimal routes.</li> </ul>                                                                                                                                                                                      |
| Original nexthop | Original next hop of the route. If the route was obtained from a BGP update message, the original next hop is the next hop IP address in the message.                                                                                                                                                                                                |
| OutLabel         | Outgoing label of the route.                                                                                                                                                                                                                                                                                                                         |
| RxPathID         | Add-Path ID of received routes.                                                                                                                                                                                                                                                                                                                      |
| TxPathID         | Add-Path ID of advertised routes.                                                                                                                                                                                                                                                                                                                    |
| AS-path          | AS_PATH attribute of the route, which records the ASs the route has passed and avoids routing loops.                                                                                                                                                                                                                                                 |
| Origin           | Origin of the route: <ul style="list-style-type: none"> <li>• <b>igp</b>—Originated in the AS. The origin of routes advertised with the <b>network</b> command is IGP.</li> <li>• <b>egp</b>—Learned through EGP.</li> <li>• <b>incomplete</b>—Unknown origin. The origin of routes redistributed from IGP protocols is INCOMPLETE.</li> </ul>       |
| Attribute value  | BGP path attributes: <ul style="list-style-type: none"> <li>• <b>MED</b>—MED value.</li> <li>• <b>localpref</b>—Local preference value.</li> <li>• <b>pref-val</b>—Preferred value.</li> <li>• <b>pre</b>—Route preference.</li> </ul>                                                                                                               |
| State            | Current state of the route: <ul style="list-style-type: none"> <li>• valid.</li> <li>• internal.</li> <li>• external.</li> <li>• local.</li> <li>• synchronize.</li> <li>• best.</li> <li>• delay. (The value means optimal route selection will be delayed for this route. This field is displayed only in detailed output information.)</li> </ul> |
| From             | IP address of the BGP peer that advertised the route.                                                                                                                                                                                                                                                                                                |
| Rely Nexthop     | Next hop found by route recursion. If no next hop is found, this field displays <b>not resolved</b> .                                                                                                                                                                                                                                                |
| IP precedence    | IP precedence in the range of 0 to 7. <b>N/A</b> indicates that the route does not support this field.                                                                                                                                                                                                                                               |
| QoS local ID     | QoS local ID in the range of 1 to 4095. <b>N/A</b> indicates that the route does not support this field.                                                                                                                                                                                                                                             |
| Traffic index    | Traffic index in the range of 1 to 64. <b>N/A</b> indicates that the route does not support this field.                                                                                                                                                                                                                                              |

## Modified command: display bgp routing-table ipv6 multicast

### Syntax

```
display bgp [instance instance-name] routing-table ipv6 multicast
[ipv6-address prefix-length [advertise-info] | as-path-acl
as-path-acl-number | community-list { { basic-community-list-number |
comm-list-name } [whole-match] | adv-community-list-number } | peer
ipv6-address { advertised-routes | received-routes } [ipv6-address
prefix-length | statistics] | statistics]
```

### Views

Any view

### Change description

After modification, the **delay** route state can be displayed in the command output, as shown in the following example:

# Display detailed information about BGP IPv6 multicast routes destined to network 2::2/128.

```
<Sysname> display bgp routing-table ipv6 multicast 2::2 128
```

```
BGP local router ID: 192.168.1.139
```

```
Local AS number: 100
```

```
Paths: 1 available, 1 best
```

```
BGP routing table information of 2::2/128:
```

```
Imported route.
```

```
Original nexthop: ::1
```

```
OutLabel : NULL
```

```
RxPathID : 0x0
```

```
TxPathID : 0x0
```

```
AS-path : (null)
```

```
Origin : incomplete
```

```
Attribute value : MED 0, pref-val 32768
```

```
State : valid, local, best, delay
```

```
IP precedence : N/A
```

```
QoS local ID : N/A
```

```
Traffic index : N/A
```

**Table 63 Command output**

| Field            | Description                                                                                                                                              |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Paths            | Number of routes: <ul style="list-style-type: none"><li><b>available</b>—Number of valid routes.</li><li><b>best</b>—Number of optimal routes.</li></ul> |
| Original nexthop | Original next hop of the route. If the route was obtained from a BGP update message, the original next hop is the next hop IP address in the message.    |
| OutLabel         | Outgoing label of the route.                                                                                                                             |
| RxPathID         | Add-Path ID of received routes.                                                                                                                          |
| TxPathID         | Add-Path ID of advertised routes.                                                                                                                        |

| Field           | Description                                                                                                                                                                                                                                                                                                                                          |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AS-path         | AS_PATH attribute of the route, which records the ASs the route has passed and avoids routing loops.                                                                                                                                                                                                                                                 |
| Origin          | Origin of the route: <ul style="list-style-type: none"> <li>• <b>igp</b>—Originated in the AS. The origin of routes advertised with the <b>network</b> command is IGP.</li> <li>• <b>egp</b>—Learned through EGP.</li> <li>• <b>incomplete</b>—Unknown origin. The origin of routes redistributed from IGP protocols is INCOMPLETE.</li> </ul>       |
| Attribute value | BGP path attributes: <ul style="list-style-type: none"> <li>• <b>MED</b>—MED value.</li> <li>• <b>localpref</b>—Local preference value.</li> <li>• <b>pref-val</b>—Preferred value.</li> <li>• <b>pre</b>—Route preference.</li> </ul>                                                                                                               |
| State           | Current state of the route: <ul style="list-style-type: none"> <li>• valid.</li> <li>• internal.</li> <li>• external.</li> <li>• local.</li> <li>• synchronize.</li> <li>• best.</li> <li>• delay. (The value means optimal route selection will be delayed for this route. This field is displayed only in detailed output information.)</li> </ul> |
| From            | IP address of the BGP peer that advertised the route.                                                                                                                                                                                                                                                                                                |
| Rely Nexthop    | Next hop found by route recursion. If no next hop is found, this field displays <b>not resolved</b> .                                                                                                                                                                                                                                                |
| IP precedence   | IP precedence in the range of 0 to 7. <b>N/A</b> indicates that the route does not support this field.                                                                                                                                                                                                                                               |
| QoS local ID    | QoS local ID in the range of 1 to 4095. <b>N/A</b> indicates that the route does not support this field.                                                                                                                                                                                                                                             |
| Traffic index   | Traffic index in the range of 1 to 64. <b>N/A</b> indicates that the route does not support this field.                                                                                                                                                                                                                                              |

Modified command: display bgp routing-table ipv6 unicast

### Syntax

```
display bgp [instance instance-name] routing-table ipv6 [unicast]
[vpn-instance vpn-instance-name] [ipv6-address prefix-length
[advertise-info] | as-path-acl as-path-acl-number | community-list
{ { basic-community-list-number | comm-list-name } [whole-match] |
adv-community-list-number } | peer ipv6-address { advertised-routes |
received-routes } [ipv6-address prefix-length | statistics] |
statistics]
```

### Views

Any view

## Change description

After modification, the **delay** route state can be displayed in the command output, as shown in the following example:

# Display detailed information about BGP IPv6 unicast routes destined to network 2::/64.

```
<Sysname> display bgp routing-table ipv6 2:: 64
```

```
BGP local router ID: 192.168.1.135
```

```
Local AS number: 200
```

```
Paths: 2 available, 1 best
```

```
BGP routing table information of 2::/64:
```

```
From : 10.1.1.1 (192.168.1.136)
```

```
Rely nexthop : ::FFFF:10.1.1.1
```

```
Original nexthop: ::FFFF:10.1.1.1
```

```
OutLabel : NULL
```

```
RxPathID : 0x0
```

```
TxPathID : 0x0
```

```
AS-path : 100
```

```
Origin : igp
```

```
Attribute value : MED 0, pref-val 0
```

```
State : valid, external, best, delay
```

```
IP precedence : N/A
```

```
QoS local ID : N/A
```

```
Traffic index : N/A
```

```
Backup route.
```

```
From : 1::1 (192.168.1.136)
```

```
Rely nexthop : 1::1
```

```
Original nexthop: 1::1
```

```
OutLabel : NULL
```

```
RxPathID : 0x0
```

```
TxPathID : 0x0
```

```
AS-path : 100
```

```
Origin : igp
```

```
Attribute value : MED 0, pref-val 0
```

```
State : valid, external
```

```
IP precedence : N/A
```

```
QoS local ID : N/A
```

```
Traffic index : N/A
```

**Table 64 Command output**

| Field | Description                                                                                                                                                  |
|-------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Paths | Number of routes: <ul style="list-style-type: none"><li>• <b>available</b>—Number of valid routes.</li><li>• <b>best</b>—Number of optimal routes.</li></ul> |

| Field            | Description                                                                                                                                                                                                                                                                                                                                    |
|------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Original nexthop | Original next hop of the route. If the route was obtained from a BGP update message, the original next hop is the next hop IP address in the message.                                                                                                                                                                                          |
| OutLabel         | Outgoing label of the route.                                                                                                                                                                                                                                                                                                                   |
| RxPathID         | Add-Path ID of received routes.                                                                                                                                                                                                                                                                                                                |
| TxPathID         | Add-Path ID of advertised routes.                                                                                                                                                                                                                                                                                                              |
| AS-path          | AS_PATH attribute of the route, which records the ASs the route has passed and avoids routing loops.                                                                                                                                                                                                                                           |
| Origin           | Origin of the route: <ul style="list-style-type: none"> <li>• <b>igp</b>—Originated in the AS. The origin of routes advertised with the <b>network</b> command is IGP.</li> <li>• <b>egp</b>—Learned through EGP.</li> <li>• <b>incomplete</b>—Unknown origin. The origin of routes redistributed from IGP protocols is INCOMPLETE.</li> </ul> |
| Attribute value  | BGP path attributes: <ul style="list-style-type: none"> <li>• <b>MED</b>—MED value.</li> <li>• <b>localpref</b>—Local preference value.</li> <li>• <b>pref-val</b>—Preferred value.</li> <li>• <b>pre</b>—Route preference.</li> </ul>                                                                                                         |
| State            | Current state of the route: <ul style="list-style-type: none"> <li>• valid.</li> <li>• internal.</li> <li>• external.</li> <li>• local.</li> <li>• best.</li> <li>• delay. (The value means optimal route selection will be delayed for this route. This field is displayed only in detailed output information.)</li> </ul>                   |
| From             | IP address of the BGP peer that advertised the route.                                                                                                                                                                                                                                                                                          |
| Rely Nexthop     | Next hop found by route recursion. If no next hop is found, this field displays <b>not resolved</b> .                                                                                                                                                                                                                                          |
| IP precedence    | IP precedence in the range of 0 to 7. <b>N/A</b> indicates that the route does not support this field.                                                                                                                                                                                                                                         |
| QoS local ID     | QoS local ID in the range of 1 to 4095. <b>N/A</b> indicates that the route does not support this field.                                                                                                                                                                                                                                       |
| Traffic index    | Traffic index in the range of 1 to 64. <b>N/A</b> indicates that the route does not support this field.                                                                                                                                                                                                                                        |

## Modified feature: Applying a routing policy to routes outgoing to a peer or peer group

### Feature change description

From this release, you can use the **apply as-path** command to modify the AS\_PATH attribute in the routing policy applied to routes outgoing to a peer or peer group.

## Command changes

None.

## Modified feature: Specifying an ACL

### Feature change description

You can specify a Layer 2 ACL or user-defined ACL by only its number in the following commands, without specifying the **mac** or **user-defined** keyword:

- `acl copy`
- `display packet-filter statistics`
- `display packet-filter statistics sum`
- `display packet-filter verbose`
- `packet-filter` (Ethernet service instance view)
- `packet-filter` (interface view)
- `packet-filter vlan-interface`
- `reset acl counter`
- `reset packet-filter statistics`

## Command changes

None.

## Modified feature: Defining an ACL match criterion in a traffic class of a QoS policy

### Feature change description

You can use a Layer 2 ACL or user-defined ACL as an ACL match criterion in a traffic class of a QoS policy.

## Command changes

### Modified command: if-match

#### Syntax

```
if-match match-criteria
```

#### Views

Traffic class view

#### Change description

Before modification: Only basic ACLs and advanced ACLs can be used as an ACL match criterion.

After modification: Layer 2 ACL or user-defined ACL can also be used as an ACL match criterion.

# Modified feature: Applying a QoS policy globally

## Feature change description

You can apply accounting-type QoS policies and marking-type QoS policies globally.

## Command changes

### Modified command: qos apply policy global

#### Old syntax

```
qos apply policy policy-name global { inbound | outbound }
```

#### New syntax

```
qos apply [accounting | remarking] policy policy-name global { inbound
| outbound }
```

#### Views

System view

#### Parameters

**accounting:** Applies an accounting-type QoS policy.

**marking:** Applies a marking-type QoS policy.

#### Change description

Before modification: You can apply only generic QoS policies globally.

After modification: You can also apply accounting-type QoS policies and marking-type QoS policies globally.

### Modified command: display qos policy global

#### Old syntax

```
display qos policy global [slot slot-number] [inbound | outbound]
```

#### New syntax

```
display qos [accounting | remarking] policy global [slot slot-number]
[inbound | outbound]
```

#### Views

Any view

#### Parameters

**accounting:** Displays information about accounting-type QoS policies.

**marking:** Displays information about marking-type QoS policies.

#### Change description

Before modification: You can display information about only generic QoS policies applied globally.

After modification: You can also display information about accounting-type QoS policies and marking-type QoS policies applied globally.

## Modified command: display qos policy diagnosis global

### Old syntax

```
display qos policy diagnosis global [slot slot-number] [inbound |
outbound]
```

### New syntax

```
display qos [accounting | remarking] policy diagnosis global [slot
slot-number] [inbound | outbound]
```

### Views

Any view

### Parameters

**accounting:** Displays diagnostic information about accounting-type QoS policies.

**marking:** Displays diagnostic information about marking-type QoS policies.

### Change description

Before modification: You can display diagnostic information about only generic QoS policies applied globally.

After modification: You can also display diagnostic information about accounting-type QoS policies and marking-type QoS policies applied globally.

## Modified command: reset qos policy global

### Old syntax

```
reset qos policy global [inbound / outbound]
```

### New syntax

```
reset qos [accounting | remarking] policy global [inbound / outbound]
```

### Views

Any view

### Parameters

**accounting:** Clears the statistics for accounting-type QoS policies.

**marking:** Clears the statistics for marking-type QoS policies.

### Change description

Before modification: You can clear the statistics for only generic QoS policies applied globally.

After modification: You can also clear the statistics for accounting-type QoS policies and marking-type QoS policies applied globally.

## Modified feature: Configuring a test profile for RADIUS server status detection

### Feature change description

This feature added support for EAP-based detection. For a RADIUS server, the device simulates an EAP authentication with the username and password specified in the test profile used by the server.

The simulated EAP authentication starts at the beginning of each detection interval. If the EAP authentication completes within a detection interval, the device determines that the RADIUS server is available.

Simulating a complete EAP authentication process, EAP-based detection provides more reliable detection results than simple detection. As a best practice, configure EAP-based detection on a network environment where EAP authentication is configured.

If you specify a nonexistent EAP profile in a test profile, the device performs simple detection for the RADIUS servers that use the test profile. After the EAP profile is configured, the device will start EAP-based detection at the next detection interval.

## Command changes

### Modified command: radius-server test-profile

#### Old syntax

```
radius-server test-profile profile-name username name [password { cipher |
simple } string] [interval interval]
undo radius-server test-profile profile-name
```

#### New syntax

```
radius-server test-profile profile-name username name [password { cipher |
simple } string] [interval interval] [eap-profile eap-profile-name]
undo radius-server test-profile profile-name
```

#### Views

System view

#### Change description

Before modification: No parameter is available for EAP-based RADIUS server status detection.

After modification: The **eap-profile** *eap-profile-name* option was added to this command. This option specifies an EAP profile for EAP-based RADIUS server status detection.

## Modified feature: RADIUS server quiet timer

### Feature change description

The minimum value for the RADIUS server quiet timer was changed from 1 minute to 0 minutes.

## Command changes

### Modified command: timer quiet (RADIUS scheme view)

#### Syntax

```
timer quiet minutes
undo timer quiet
```

#### Views

RADIUS scheme view

## Change description

Before modification: The value range for the *minutes* argument is 1 to 255 minutes.

After modification: The value range for the *minutes* argument is 0 to 255 minutes. If you set this argument to 0, the device does not change the state of the current server for a user when the server is unreachable. It sends an authentication or accounting request of the user to the next server in active state. For an authentication or accounting request of a new user, it still tries to send the request to the current server because the current server is in active state.

## Modified feature: Specifying the source IP address for outgoing RADIUS packets

### Feature change description

This version added support for specifying a source interface to provide the source IP address for outgoing RADIUS packets.

### Command changes

#### Modified command: nas-ip (RADIUS scheme view)

##### Old syntax

```
nas-ip { ipv4-address | ipv6 ipv6-address }
undo nas-ip [ipv6]
```

##### New syntax

```
nas-ip { ipv4-address | interface interface-type interface-number | ipv6
ipv6-address }
undo nas-ip [interface | ipv6]
```

##### Views

RADIUS scheme view

##### Parameters

**interface interface-type interface-number**: Specifies a source interface by its type and number. The device uses the primary IPv4 address or the IPv6 address of the interface as the source IP address of an outgoing RADIUS packet. You can specify only one source interface to provide the source IP address for outgoing RADIUS packets. Make sure the route between the source interface and the RADIUS server is reachable.

## Change description

Before modification: You cannot specify a source interface to provide the source IP address for outgoing RADIUS packets.

After modification: The **interface interface-type interface-number** option was added to this command. You can use this option to specify a source interface to provide the source IP address for outgoing RADIUS packets.

## Modified command: radius nas-ip

### Old syntax

```
radius nas-ip { ipv4-address | ipv6 ipv6-address } [vpn-instance
vpn-instance-name] }
```

```
undo radius nas-ip { ipv4-address | ipv6 ipv6-address } [vpn-instance
vpn-instance-name] }
```

### New syntax

```
radius nas-ip { interface interface-type interface-number |
{ ipv4-address | ipv6 ipv6-address } [vpn-instance vpn-instance-name] }
```

```
undo radius nas-ip { interface | { ipv4-address | ipv6 ipv6-address }
[vpn-instance vpn-instance-name] }
```

### Views

System view

### Parameters

**interface interface-type interface-number**: Specifies a source interface by its type and number. The device uses the primary IPv4 address or the IPv6 address of the interface as the source IP address of an outgoing RADIUS packet.

### Change description

Before modification: You cannot specify a source interface to provide the source IP address for outgoing RADIUS packets.

After modification: The **interface interface-type interface-number** option was added to this command. You can use this option to specify a source interface to provide the source IP address for outgoing RADIUS packets.

## Modified feature: Specifying the source IP address for outgoing HWTACACS packets

### Feature change description

This version added support for specifying a source interface to provide the source IP address for outgoing HWTACACS packets.

### Command changes

## Modified command: nas-ip (HWTACACS scheme view)

### Old syntax

```
nas-ip { ipv4-address | ipv6 ipv6-address }
```

```
undo nas-ip [ipv6]
```

### New syntax

```
nas-ip { ipv4-address | interface interface-type interface-number | ipv6
ipv6-address }
```

```
undo nas-ip [interface | ipv6]
```

## Views

HWTACACS scheme view

## Parameters

**interface** *interface-type interface-number*: Specifies a source interface by its type and number. The device uses the primary IPv4 address or the IPv6 address of the interface as the source IP address of an outgoing HWTACACS packet. You can specify only one source interface to provide the source IP address for outgoing HWTACACS packets. Make sure the route between the source interface and the HWTACACS server is reachable.

## Change description

Before modification: You cannot specify a source interface to provide the source IP address for outgoing HWTACACS packets.

After modification: The **interface** *interface-type interface-number* option was added to this command. You can use this option to specify a source interface to provide the source IP address for outgoing HWTACACS packets.

## Modified command: hwtacacs nas-ip

### Old syntax

```
hwtacacs nas-ip { ipv4-address | ipv6 ipv6-address } [vpn-instance vpn-instance-name] }
undo hwtacacs nas-ip { ipv4-address | ipv6 ipv6-address } [vpn-instance vpn-instance-name] }
```

### New syntax

```
hwtacacs nas-ip { interface interface-type interface-number |
 { ipv4-address | ipv6 ipv6-address } [vpn-instance vpn-instance-name] }
undo hwtacacs nas-ip { interface | { ipv4-address | ipv6 ipv6-address }
 [vpn-instance vpn-instance-name] }
```

## Views

System view

## Parameters

**interface** *interface-type interface-number*: Specifies a source interface by its type and number. The device uses the primary IPv4 address or the IPv6 address of the interface as the source IP address of an outgoing HWTACACS packet.

## Change description

Before modification: You cannot specify a source interface to provide the source IP address for outgoing HWTACACS packets.

After modification: The **interface** *interface-type interface-number* option was added to this command. You can use this option to specify a source interface to provide the source IP address for outgoing HWTACACS packets.

# Modified feature: Including user IP addresses in MAC authentication requests

## Feature change description

This feature added support for specifying an ACL-based filter to identify source IP addresses that can or cannot trigger MAC authentication.

If the user host is configured with IPv6, the device might receive packets that contain an IPv6 link-local address, which starts with fe80. MAC authentication failure will occur if this address is used in MAC authentication. To avoid MAC authentication failure, configure a basic ACL to exclude the IPv6 IP addresses that start with fe80.

When you configure the ACL, follow these guidelines:

- Use permit rules to identify source IP addresses that are valid for MAC authentication. Use deny rules to identify source IP addresses that cannot trigger MAC authentication.
- In the rules, only the action keyword (permit or deny) and the source IP match criterion can take effect.
- As a best practice, configure a deny rule to exclude the IPv6 IP addresses that start with fe80 from triggering MAC authentication.
- If you configure permit rules, add a **deny all** rule at the bottom of the ACL.

## Command changes

### Modified command: mac-authentication carry user-ip

#### Old syntax

```
mac-authentication carry user-ip
```

#### New syntax

```
mac-authentication carry user-ip [exclude-ip acl acl-number]
```

#### Views

Layer 2 Ethernet interface view

#### Parameters

**exclude-ip**: Specifies an ACL-based filter to identify source IP addresses that can or cannot trigger MAC authentication.

**acl acl-number**: Specifies a basic ACL. The value range for the *acl-number* argument is 2000 to 2999.

#### Change description

Before modification: This feature does not support specifying an ACL-based filter to identify source IP addresses that can or cannot trigger MAC authentication.

After modification: The **exclude-ip acl acl-number** option was added to this command for specifying an ACL-based filter to identify source IP addresses that can or cannot trigger MAC authentication.

# Modified feature: Configuring MAC-based MAC authentication user accounts

## Feature change description

Support for password configuration was added to MAC-based MAC authentication user accounts.

## Command changes

### Modified command: mac-authentication user-name-format

#### Old syntax

```
mac-authentication user-name-format { fixed [account name] [password
{ cipher | simple } string] | mac-address [{ with-hyphen | without-hyphen }
[lowercase | uppercase]] }
```

#### New syntax

```
mac-authentication user-name-format { fixed [account name] | mac-address
[{ with-hyphen | without-hyphen } [lowercase | uppercase]] } [password
{ cipher | simple } string]
```

## Views

System view

## Change description

Before modification: You cannot specify a password for MAC-based MAC authentication user accounts. The MAC address of each user is used as the password.

After modification: You can specify a password for all MAC-based MAC authentication user accounts by using the **password { cipher | simple } *string*** option. If you do not specify a password, each user uses its own MAC address as the password.

- **password:** Specifies the password for MAC-based MAC authentication user accounts.
- **cipher:** Specifies a password in encrypted form.
- **simple:** Specifies a password in plaintext form. For security purposes, the password specified in plaintext form will be stored in encrypted form.
- ***string*:** Specifies the password. Its plaintext form is a case-sensitive string of 1 to 63 characters. Its encrypted form is a case-sensitive string of 1 to 117 characters.

# Modified feature: Handling new MAC access attempts in a VLAN after port security's MAC address limit for that VLAN is reached

## Feature change description

This version added support for logging the first access attempt from a new MAC address in a VLAN after port security's MAC address limit for that VLAN is reached. For each VLAN, the system does not log any access attempts from a new MAC address except the first one after the MAC address limit is reached.

## Command changes

### Old syntax

```
port-security access-user log enable [failed-authorization |
mac-learning | violation] *

undo port-security access-user log enable [failed-authorization |
mac-learning | violation] *
```

### New syntax

```
port-security access-user log enable [failed-authorization |
mac-learning | violation | vlan-mac-limit] *

undo port-security access-user log enable [failed-authorization |
mac-learning | violation | vlan-mac-limit] *
```

### Views

System view

### Change description

Before modification: The system does not log any access attempts from new MAC accesses in a VLAN after port security's MAC address limit for that VLAN is reached.

After modification: The **vlan-mac-limit** keyword was added to this command. Specify this keyword to log the first access attempt from a new MAC access in a VLAN after port security's MAC address limit for that VLAN is reached.

## Modified feature: Port security NTK feature

### Feature change description

In this release, the ntkauto mode was added to the need to know (NTK) feature of port security. This feature specifies the ntkauto mode. A port in ntkauto mode forwards broadcast frames, multicast frames, and unicast frames with authenticated destination MAC addresses only when the port has online users.

## Command changes

### Modified command: port-security ntk-mode

#### Old syntax

```
port-security ntk-mode { ntk-withbroadcasts | ntk-withmulticasts | ntkonly }
```

#### New syntax

```
port-security ntk-mode { ntk-withbroadcasts | ntk-withmulticasts |
ntkauto | ntkonly }
```

### Views

Layer 2 Ethernet interface view

### Change description

Before modification: The **ntkauto** keyword was not added to this command.

After modification: The **ntkauto** keyword was not added to this command.

# Modified feature: Password handling manners with password control enabled globally

## Feature change description

The device changed the password handling manners for device management users and user roles when password control is enabled globally.

## Password handling manners for device management users

### Before modification

When password control is enabled globally, the device handles passwords of device management users in the following manners:

- In the history records, passwords specified in plaintext form are saved in encrypted form. Passwords specified in hashed form are not saved in the history records.
- If a user changes its own password or the password of another user in plaintext form, the system compares the new password with the following passwords:
  - All passwords of the target user in the history records.
  - The current password of the target user.

The new password must have a minimum of four characters different from the target user's all passwords in the history records and current password.

- If a user changes its own password or the password of another user in hashed form, the system does not compare the new password with passwords in the history records and the current password.
- If a user deletes its own password, the system does not request the user to enter the current password in plaintext form.
- In FIPS mode, if a network-admin user changes its own password, the system does not request the user to enter the current password in plaintext form.

### After modification

When password control is enabled globally, the device handles passwords of device management users in the following manners:

- All passwords in the history records are saved in hashed form.
- If a user changes its own password in plaintext form, the system requests the user to enter the current password in plaintext form. The new password must be different from all passwords of the user in the history records and the current password. In addition, the new password must have a minimum of four characters different from the current password.
- If a user changes the password for another user in plaintext form, the new password must be different from the target user's all passwords in the history records and current password.
- If a user changes its own password or the password of another user in hashed form, the system does not compare the new password with passwords in the history records and the current password.
- If a user deletes its own password, the system requests the user to enter the current password in plaintext form.
- In FIPS mode, if a network-admin user changes its own password, the system requests the user to enter the current password in plaintext form.

# Password handling manners for obtaining temporary user role authorization

## Before modification

When password control is enabled globally, the device handles passwords for obtaining temporary user role authorization in the following manners:

- In the history records, passwords specified in plaintext form are saved in encrypted form. Passwords specified in hashed form are not saved in the history records.
- If you change the password for a user role in plaintext form, the new password must be different from all passwords in the history records and the current password of the user role. In addition, the new password must have a minimum of four characters different from all passwords in the history records and the current password.
- If you change the password for a user role in hashed form, the system does not compare the new password with passwords in the history records and the current password.

## After modification

When password control is enabled globally, the device handles passwords for obtaining temporary user role authorization in the following manners:

- All passwords in the history records are saved in hashed form.
- If you change the password for a user role in plaintext form, the new password must be different from all passwords in the history records and the current password of the user role.
- If you change the password for a user role in hashed form, the system does not compare the new password with passwords in the history records and the current password.

## Command changes

None.

# Modified feature: RSA key modulus length

## Feature change description

The supported RSA key modulus length was changed.

## Command changes

### Modified command: public-key local create

#### Syntax

In non-FIPS mode:

```
public-key local create { dsa | ecdsa [secp192r1 | secp256r1 | secp384r1
| secp521r1] | rsa } [name key-name]
```

In FIPS mode:

```
public-key local create { dsa | ecdsa [secp256r1 | secp384r1 | secp521r1]
| rsa } [name key-name]
```

#### Views

System view

## Change description

Before modification, the key modulus length of a local RSA key pair is as follows:

- In non-FIPS mode:
  - Value range: 512 to 2048 bits.
  - Default: 1024 bits.To ensure security, use a minimum of 768 bits.
- In FIPS mode: 2048 bits.

After modification, the key modulus length of a local RSA key pair is as follows:

- In non-FIPS mode:
  - Value range: 512 to 4096 bits.
  - Default: 1024 bits.To ensure security, use a minimum of 768 bits.
- In FIPS mode:
  - Value range: A multiple of 256 bits in the range of 2048 to 4096 bits.
  - Default: 2048 bits.

## Modified feature: Key modulus length of the RSA key pair used for certificate request in a PKI domain

### Feature change description

The supported key modulus length of the RSA key pair used for certificate request in a PKI domain was changed.

### Command changes

Modified command: `public-key rsa`

#### Syntax

```
public-key rsa { { encryption name encryption-key-name [length key-length]
| signature name signature-key-name [length key-length] } * | general name
key-name [length key-length] }
```

#### Views

PKI domain view

## Change description

Before modification, the value range and default setting for the *key-length* argument is as follows:

- In non-FIPS mode:
  - Value range: 512 to 2048 bits.
  - Default: 1024 bits.
- In FIPS mode: 2048 bits.

After modification, the value range and default setting for the *key-length* argument is as follows:

- In non-FIPS mode:

- Value range: 512 to 4096 bits.
- Default: 1024 bits.
- In FIPS mode:
  - Value range: A multiple of 256 bits in the range of 2048 to 4096 bits.
  - Default: 2048 bits.

The longer the key modulus length, the higher the security, and the longer the key generation time.

## Modified feature: Configuring the ECDSA signature authentication method in an IKE proposal

### Feature change description

From this version, you can specify the ECDSA signature authentication method in an IKE proposal. For successful IKE negotiation with a peer by using the ECDSA signature authentication method, the peer must have a CA-signed identity certificate.

### Command changes

#### Modified command: authentication-method

##### Old syntax

```
authentication-method { dsa-signature | pre-share | rsa-signature }
undo authentication-method
```

##### New syntax

```
authentication-method { dsa-signature | ecdsa-signature | pre-share |
rsa-signature }
undo authentication-method
```

##### Views

IKE proposal view

##### Change description

The **ecdsa-signature** keyword was added to the **authentication-method** command.

## Modified feature: Specifying the DH group used for key negotiation in IKE phase 1

### Feature change description

From this version, you can specify the 256-bit or 384-bit Diffie-Hellman group in an IKE proposal for key negotiation in IKE phase 1.

## Command changes

### Modified command: `dh`

#### Old syntax

In non-FIPS mode:

```
dh { group1 | group14 | group2 | group24 | group5 }
undo dh
```

In FIPS mode:

```
dh group14
undo dh
```

#### New syntax

In non-FIPS mode:

```
dh { group1 | group14 | group19 | group2 | group20 | group24 | group5 }
undo dh
```

In FIPS mode:

```
dh group14
undo dh
```

#### Views

IKE proposal view

#### Change description

The following keywords were added to the `dh` command in non-FIPS mode:

- `group19`: Uses the 256-bit Diffie-Hellman group.
- `group20`: Uses the 384-bit Diffie-Hellman group.

## Modified feature: Keyboard-interactive authentication support for SSH users

### Feature change description

In this release, the keyboard-interactive authentication method is supported when you create an SSH user on the device that acts as an SSH server.

The keyboard-interactive authentication method supports multiple rounds of interactive exchanges of information. To pass authentication, the user must provide all interactive information required by the remote authentication server. If the remote authentication server does not require interactive information, the keyboard-interactive authentication process is the same as the password authentication.

## Command changes

### Modified command: ssh user

#### Old syntax

In non-FIPS mode:

```
ssh user username service-type { all | netconf | scp | sftp | stelnet }
authentication-type { password | { any | password-publickey | publickey }
[assign { pki-domain domain-name | publickey keyname&<1-6> }] }
```

```
undo ssh user username
```

In FIPS mode:

```
ssh user username service-type { all | netconf | scp | sftp | stelnet }
authentication-type { password | password-publickey [assign { pki-domain
domain-name | publickey keyname&<1-6> }] }
```

```
undo ssh user username
```

#### New syntax

In non-FIPS mode:

```
ssh user username service-type { all | netconf | scp | sftp | stelnet }
authentication-type { keyboard-interactive | password | { any |
password-publickey | publickey } [assign { pki-domain domain-name | publickey
keyname&<1-6> }] }
```

```
undo ssh user username
```

In FIPS mode:

```
ssh user username service-type { all | netconf | scp | sftp | stelnet }
authentication-type { keyboard-interactive | password | password-publickey
[assign { pki-domain domain-name | publickey keyname&<1-6> }] }
```

```
undo ssh user username
```

#### Views

System view

#### Change description

Before modification: The keyboard-interactive authentication method is not supported when you create an SSH user on the device that acts as an SSH server.

After modification: The keyboard-interactive authentication method is supported when you create an SSH user on the device that acts as an SSH server.

## Modified feature: Displaying IPv4 source guard bindings

### Feature change description

From this version, the **arp-snooping-vsi** keyword is added to the **display ip source binding** command, the **arp-suppression** keyword is deleted, and the **arp-snooping** keyword is changed to **arp-snooping-vlan**.

## Command changes

### Old syntax

```
display ip source binding [static | [vpn-instance vpn-instance-name]
[arp-snooping | arp-suppression | dhcp-relay | dhcp-server | dhcp-snooping
| dot1x]] [ip-address ip-address] [mac-address mac-address] [vlan
vlan-id] [interface interface-type interface-number] [slot
slot-number]
```

### New syntax

```
display ip source binding [static | [vpn-instance vpn-instance-name]
[arp-snooping-vlan | arp-snooping-vsi | dhcp-relay | dhcp-server |
dhcp-snooping | dot1x]] [ip-address ip-address] [mac-address
mac-address] [vlan vlan-id] [interface interface-type interface-number]
[slot slot-number]
```

### Views

Any view

### Parameters

**arp-snooping-vlan**: Specifies IPv4SG bindings generated based on ARP snooping for VLANs.

**arp-snooping-vsi**: Specifies IPv4SG bindings generated based on ARP snooping for VSIs.

### Change description

Before modification: The **arp-snooping-vsi** keyword is not supported and the **arp-suppression** keyword is supported.

After modification: The keyword change is as follows:

- The **arp-snooping-vsi** keyword is supported.
- The **arp-suppression** is deleted.
- The **arp-snooping** keyword is changed to **arp-snooping-vlan**.

## Modified feature: Displaying IPv6 source guard bindings

### Feature change description

From this version, the **nd-snooping-vsi** keyword is added to the **display ipv6 source binding** command, the **nd-suppression** keyword is deleted, and the **nd-snooping** keyword is changed to **nd-snooping-vlan**.

## Command changes

### Old syntax

```
display ipv6 source binding [static | [vpn-instance vpn-instance-name]
[nd-snooping | nd-suppression | dhcpv6-relay | dhcpv6-snooping | dot1x]]
[ip-address ipv6-address] [mac-address mac-address] [vlan vlan-id]
[interface interface-type interface-number] [slot slot-number]
```

### New syntax

```
display ipv6 source binding [static | [vpn-instance vpn-instance-name]
[nd-snooping-vlan | nd-snooping-vsi | dhcpv6-relay | dhcpv6-snooping |
```

```
dot1x]] [ip-address ipv6-address] [mac-address mac-address] [vlan
vlan-id] [interface interface-type interface-number] [slot
slot-number]
```

## Views

Any view

## Parameters

**nd-snooping-vlan**: Specifies IPv6SG bindings generated based on ND snooping for VLANs.

**nd-snooping-vsi**: Specifies IPv6SG bindings generated based on ND snooping for VSIs.

## Change description

Before modification: The **nd-snooping-vsi** keyword is not supported and the **nd-suppression** keyword is supported.

After modification: The keyword change is as follows:

- The **nd-snooping-vsi** keyword is supported.
- The **nd-suppression** is deleted.
- The **nd-snooping** keyword is changed to **nd-snooping-vlan**.

# Modified feature: Configuring ARP attack detection logging

## Feature change description

Support for specifying the maximum number of ARP attack detection log entries during each log output was added.

## Command changes

### Modified command: arp detection log enable

#### Old syntax

```
arp detection log enable [interval interval]
```

#### New syntax

```
arp detection log enable [interval interval | number number]
```

## Views

System view

## Change description

Before modification: The **number** keyword was not supported.

After modification: Support for the **number** *number* option was added. You can specify the maximum number of ARP attack detection log entries for each log output.

## Modified feature: Creating a BFD session for detecting the local interface state

### Feature change description

From this release, you can specify a parameter template when creating a BFD session for detecting the local interface state.

### Command changes

#### Modified command: bfd detect-interface

##### Old syntax

```
bfd detect-interface source-ip ip-address [discriminator local
local-value remote remote-value]
```

##### New syntax

```
bfd detect-interface source-ip ip-address [discriminator local
local-value remote remote-value] [template template-name]
```

##### Views

Interface view

##### Change description

Before modification: The **template** keyword is not supported. The BFD session uses the BFD parameters configured in interface view.

After modification: The **template** keyword is supported. If you specify this keyword, the BFD session uses the BFD parameters configured in the template. If you do not specify this keyword, the BFD session uses the BFD parameters configured in interface view.

## Modified feature: Configuring the BFD authentication mode for single-hop BFD control packets

### Feature change description

From this release, the following algorithms are supported for authenticating single-hop BFD control packets:

- HMAC MD5 algorithm.
- HMAC Meticulous MD5 algorithm.
- HMAC Meticulous SHA1 algorithm.
- HMAC SHA1 algorithm.

## Command changes

### Modified command: bfd authentication-mode

#### Old syntax

```
bfd authentication-mode { m-md5 | m-sha1 | md5 | sha1 | simple } key-id
{ cipher | plain } string
```

#### New syntax

```
bfd authentication-mode { hmac-md5 | hmac-mmd5 | hmac-msha1 | hmac-sha1
| m-md5 | m-sha1 | md5 | sha1 | simple } key-id { cipher | plain } string
```

#### Views

Interface view

BFD template view

#### Change description

Before modification: The **hmac-md5**, **hmac-mmd5**, **hmac-msha1**, and **hmac-sha1** keywords are not supported.

After modification: The **hmac-md5**, **hmac-mmd5**, **hmac-msha1**, and **hmac-sha1** keywords are supported.

## Modified feature: Setting the minimum interval for receiving BFD echo packets

### Feature change description

From this release, the configuration of the minimum interval for receiving BFD echo packets is not supported in BFD template view.

## Command changes

### Modified command: bfd min-echo-receive-interval

#### Syntax

```
bfd min-echo-receive-interval interval
```

#### Old views

Interface view

BFD template view

#### New views

Interface view

#### Change description

Before modification: The **bfd min-echo-receive-interval** command is available in both interface view and BFD template view.

After modification: The **bfd min-echo-receive-interval** command is available only in interface view.

## Modified feature: Support for specific BFD commands in VSI interface view

### Feature change description

From this release, certain BFD commands can be configured in VSI interface view:

### Command changes

#### Syntax

```
bfd authentication-mode
bfd demand enable
bfd detect-multiplier
bfd echo enable
bfd min-echo-receive-interval
bfd min-receive-interval
bfd min-transmit-interval
```

#### Change description

Before modification: The commands are not available in VSI interface view.

After modification: The commands are available in VSI interface view.

## Modified feature: Setting the maximum number of active routes in a VPN instance

### Feature change description

The maximum number of active routes supported in a VPN instance was changed.

### Command changes

#### Modified command: routing-table limit

#### Syntax

```
routing-table limit number { warn-threshold | simply-alert }
undo routing-table limit
```

#### Views

VPN instance view  
VPN instance IPv4 VPN view  
VPN instance IPv6 VPN view

## Change description

Before modification: The maximum value for the *number* argument is 262144 in VPN instance view, 393216 in VPN instance IPv4 VPN view, and 262144 in VPN instance IPv6 VPN view.

After modification: The maximum value for the *number* argument is 8192 in VPN instance view, 16384 in VPN instance IPv4 VPN view, and 8192 in VPN instance IPv6 VPN view.

## Modified feature: Enabling BFD on an MPLS TE tunnel interface

### Feature change description

From this release, BFD for MPLS TE does not support specifying a BFD session parameter template for the MPLS TE backup tunnel.

### Command changes

Modified command: `mpls bfd`

#### Old syntax

```
mpls bfd [[discriminator local local-id remote remote-id] [template
template-name] [backup-path template template-name] | echo]
```

#### New syntax

```
mpls bfd [[discriminator local local-id remote remote-id] [template
template-name] | echo]
```

#### Views

Tunnel interface view

## Change description

The `backup-path template template-name` option was removed from the command syntax.

## Modified feature: Setting the length of ICMP or ICMPv6 echo requests for the ping operation

### Feature change description

In the ping operation, the value range for the length of ICMP or ICMPv6 echo requests was changed from 20 to 8100 bytes to 20 to 9600 bytes.

### Command changes

Modified command: `ping`

#### Syntax

```
ping [ip] [-a source-ip | -c count | -f | -h ttl | -i interface-type
interface-number | -m interval | -n | -p pad | -q | -r | -s packet-size | -t
timeout | -tos tos | -v | vpn-instance vpn-instance-name] * host
```

## Views

Any view

## Parameters

**-s** *packet-size*: Specifies the length (in bytes) of ICMP echo requests (excluding the IP packet header and the ICMP packet header). The value range is 20 to 9600, and the default is 56.

## Change description

Before modification: The value range for the *packet-size* argument is 20 to 8100 bytes.

After modification: The value range for the *packet-size* argument is 20 to 9600 bytes.

## Modified command: ping ipv6

## Syntax

```
ping ipv6 [-a source-ipv6 | -c count | -i interface-type interface-number
| -m interval | -q | -s packet-size | -t timeout | -tc traffic-class | -v
| vpn-instance vpn-instance-name] * host
```

## Views

Any view

## Parameters

**-s** *packet-size*: Specifies the length (in bytes) of ICMPv6 echo requests (excluding the IPv6 packet header and the ICMPv6 packet header). The value range is 20 to 9600, and the default is 56.

## Change description

Before modification: The value range for the *packet-size* argument is 20 to 8100 bytes.

After modification: The value range for the *packet-size* argument is 20 to 9600 bytes.

# Modified feature: Specifying a source interface for NTP messages

## Feature change description

The *ip-address* argument was added to the command that specifies a source interface for NTP messages.

## Command changes

## Modified command: ntp-service source

## Old syntax

```
ntp-service source interface-type interface-number
```

## New syntax

```
ntp-service source { interface-type interface-number | ip-address }
```

## Views

System view

## Parameters

*interface-type interface-number*: Specifies an interface by its type and number.

*ip-address*: Specifies a source IP address.

## Usage guidelines

If you specify a source interface for NTP messages, the device uses the primary IP address of the specified interface as the source IP address to send NTP messages. Consequently, the destination address of the NTP response messages is the primary IP address of the source interface.

## Change description

Before modification: The command does not support specifying a source IP address. The device obtains the source IP address of NTP messages through the specified interface.

After modification: The command supports using the *ip-address* argument to specify a source IP address.

# Modified feature: PTP

## Feature change description

The SMPTE ST 2059-2 standard was added to PTP.

## Command changes

### Modified command: ptp profile

#### Old syntax

```
ptp profile { 1588v2 | 8021as }
undo ptp profile
```

#### New syntax

```
ptp profile { 1588v2 | 8021as | st2059-2 }
undo ptp profile
```

## Views

System view

## Parameter

**st2059-2**: Specifies the PTP profile as SMPTE ST 2059-2.

## Change description

Before modification: Supported PTP profiles include IEEE 1588 version 2 and IEEE 802.1AS.

After modification: Supported PTP profiles include IEEE 1588 version 2, IEEE 802.1AS, and SMPTE ST 2059-2.

### Modified command: display ptp clock

## Syntax

```
display ptp clock
```

## Views

Any view

## Change description

The SMPTE ST 2059-2 option was added to the **PTP profile** field in the output from the **display ptp clock** command.

```
<Sysname> display ptp clock
PTP profile : SMPTE ST 2059-2
PTP mode : BC
Slave only : No
Clock ID : 000FE2-FFFE-FF0000
Clock type : ToD0
 ToD direction : In
 ToD delay time : 0 (ns)
Clock domain : 0
Number of PTP ports : 2
Priority1 : 128
Priority2 : 128
Clock quality :
 Class : 248
 Accuracy : 254
 Offset (log variance) : 65535
Offset from master : -4 (ns)
Mean path delay : 101 (ns)
Steps removed : 1
Local clock time : Sun Jan 15 20:57:29 2011
```

**Table 65 Command output**

| Field       | Description                                                                                                                                                    |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PTP profile | PTP standard: <ul style="list-style-type: none"><li>• IEEE 1588 Version 2.</li><li>• IEEE 802.1AS.</li><li>• SMPTE ST 2059-2.</li></ul>                        |
| PTP mode    | Clock node type: <ul style="list-style-type: none"><li>• BC.</li><li>• E2ETC.</li><li>• E2ETC-OC.</li><li>• OC.</li><li>• P2PTC.</li><li>• P2PTC-OC.</li></ul> |
| Slave only  | Whether the OC operates in <b>Slave only</b> mode: <ul style="list-style-type: none"><li>• Yes.</li><li>• No.</li></ul>                                        |
| Clock ID    | Clock ID for the device, which uniquely identifies a PTP device.                                                                                               |
| Clock type  | Clock type of the device:                                                                                                                                      |

| Field                 | Description                                                                                                                                           |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------|
|                       | <ul style="list-style-type: none"> <li>• <b>Local</b>—Local clock.</li> <li>• <b>ToD0</b>—ToD 0 clock.</li> <li>• <b>ToD1</b>—ToD 1 clock.</li> </ul> |
| ToD direction         | Direction for the ToD clock. The value for this field is <b>In</b> . This field is not displayed when the device is a local clock.                    |
| ToD delay time        | Delay correction value in nanoseconds for the ToD clock. This field is not displayed when the device is a local clock.                                |
| Clock domain          | PTP domain where the device resides.                                                                                                                  |
| Priority1             | Priority 1 value for the device.                                                                                                                      |
| Priority2             | Priority 2 value for the device.                                                                                                                      |
| Class                 | Time class of the device.                                                                                                                             |
| Accuracy              | Time accuracy of the device.                                                                                                                          |
| Offset (log variance) | Offset of the grandmaster clock.                                                                                                                      |
| Offset from master    | Offset from the master clock node (parent clock node), in nanoseconds.<br>N/A indicates that information for this field is not obtained.              |
| Mean path delay       | Mean path delay, in nanoseconds.<br>N/A indicates that information for this field is not obtained.                                                    |
| Steps removed         | Hops from the grandmaster to the local clock node.<br>N/A indicates that information for this field is not obtained.                                  |

## Modified command: ptp announce-interval

### Syntax

**ptp announce-interval** *interval*

**undo ptp announce-interval**

### Default

If the PTP profile is IEEE 1588 version 2, the announce message sending interval is 2 ( $2^1$ ) seconds.

If the PTP profile is IEEE 802.1AS, the announce message sending interval is 1 ( $2^0$ ) second.

If the PTP profile is SMPTE ST 2059-2, the announce message sending interval is 1/4 ( $2^{-2}$ ) seconds.

### Views

Any view

### Parameters

*interval*: Specifies an exponent, to the power of which base 2 is raised to get the interval in seconds for sending announce messages. If the PTP profile is IEEE 1588 version 2, the value range for the *interval* argument is 0 to 4. If the PTP profile is IEEE 802.1AS, the value range for the *interval* argument is 0 to 6. If the PTP profile is SMPTE ST 2059-2, the value range for the *interval* argument is -3 to 1.

### Change description

The default value 1/4 ( $2^{-2}$ ) and the value range -3 to 1 were added for specifying the announce message sending interval in the SMPTE ST 2059-2 PTP profile.

## Modified command: ptp clock-source

### Syntax

```
ptp clock-source { tod0 | tod1 } { accuracy acc-value | class class-value |
time-source ts-value }

undo ptp clock-source { tod0 | tod1 } { accuracy | class | time-source }
```

### Views

System view

### Parameters

**class** *class-value*: Specifies the class of the ToD clock, in the range of 0 to 255. A smaller value indicates a higher time class. [Table 66](#) shows the description for clock classes (those not listed are reserved by the protocol and are not recommended to use).

**Table 66 Clock class description**

| Clock class (in decimal format) | Description                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 6                               | The clock node synchronizes its time to the master reference time source. PTP assigns a time table to the clock node. A clock node with time class 6 cannot become a member clock of any other clocks in the domain.                                                                                           |
| 7                               | The former time class is 6. The clock node cannot synchronize its time to a time source. It enters the reappointment mode and meets the reappointment conditions. PTP assigns a time table to the clock node. A clock node with time class 7 cannot become a member clock of any other clocks in the domain.   |
| 13                              | The clock node synchronizes its time to a time source. ARB assigns a time table to the clock node. A clock node with time class 13 cannot become a member clock of any other clocks in the domain.                                                                                                             |
| 14                              | The former time class is 13. The clock node cannot synchronize its time to a time source. It enters the reappointment mode and meets the reappointment conditions. ARB assigns a time table to the clock node. A clock node with time class 14 cannot become a member clock of any other clocks in the domain. |
| 52                              | The clock node with time class 7 becomes optional clock A because it does not meet the reappointment conditions. A clock node with time class 52 cannot become a member clock of any other clocks in the domain.                                                                                               |
| 58                              | The clock node with time class 14 becomes optional clock A because it does not meet the reappointment conditions. A clock node with time class 58 cannot become a member clock of any other clocks in the domain.                                                                                              |
| 150                             | The clock node synchronizes its frequency to a reference source with frequency accuracy of $\pm 1$ ppm (one in a million). The reference source has synchronized itself to the master reference time source. PTP assigns a time table to the clock node.                                                       |
| 158                             | The clock node synchronizes its frequency to a reference source with frequency accuracy of $\pm 10$ ppm (ten in a million). The reference source has synchronized itself to the master reference time source. PTP assigns a time table to the clock node.                                                      |
| 187                             | The clock node with time class 7 becomes optional clock B because it does not meet the reappointment conditions. A clock node with time class 187 can become a member clock of another clock in the domain.                                                                                                    |
| 193                             | The clock node with time class 14 becomes optional clock B because it does not meet the reappointment conditions. A clock node with time class 193 can become a member clock of another clock in the domain.                                                                                                   |
| 220                             | The clock node synchronizes its frequency to a reference source with frequency accuracy of $\pm 1$ ppm (one in a million). The reference source has synchronized itself to the master reference time source. ARB assigns a time table to the clock node.                                                       |

|     |                                                                                                                                                                                                                                                           |
|-----|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 228 | The clock node synchronizes its frequency to a reference source with frequency accuracy of $\pm 10$ ppm (ten in a million). The reference source has synchronized itself to the master reference time source. ARB assigns a time table to the clock node. |
| 248 | Default time class value.                                                                                                                                                                                                                                 |
| 255 | Clock node operating in <b>slave-only</b> mode.                                                                                                                                                                                                           |

**time-source** *ts-value*: Specifies the clock attribute value in the range of 0 to 255. [Table 67](#) shows the description for the clock attribute values (those not listed are reserved for the protocol).

**Table 67 Attribute value description**

| Attribute values (in hexadecimal format) | Description                                                                                                                                                                                                  |
|------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 10                                       | Atomic clock.                                                                                                                                                                                                |
| 20                                       | GPS.                                                                                                                                                                                                         |
| 30                                       | Terrestrial radio.                                                                                                                                                                                           |
| 40                                       | PTP.                                                                                                                                                                                                         |
| 50                                       | NTP.                                                                                                                                                                                                         |
| 60                                       | Handset.                                                                                                                                                                                                     |
| 90                                       | Other.                                                                                                                                                                                                       |
| A0                                       | Internal oscillator.                                                                                                                                                                                         |
| F0                                       | Grandmaster clock. The time source of the clock is a periodic sync signal. The time of the periodic sync signal is undefined and not related to real time. ARB assigns a time table to the clock node.       |
| F1                                       | Grandmaster clock. The time source of the clock is a local time source. Then clock synchronizes its time to a periodic sync signal. The clock accuracy depends on the frequency accuracy of the sync signal. |
| F2 to FE                                 | For use by alternate PTP profiles.                                                                                                                                                                           |
| FF                                       | Reserved.                                                                                                                                                                                                    |

## Change description

Clock classes 150, 158, 220, and 228 were added to the classes of the ToD clock. Clock attribute values F0 and F1 were added.

## Modified command: ptp domain

### Syntax

**ptp domain** *domain-number*

**undo ptp domain**

### Default

If the PTP profile is IEEE 1588 version 2 or IEEE 802.1AS, the PTP domain number is 0.

If the PTP profile is SMPTE ST 2059-2, the PTP domain number is 127.

### Views

System view

## Parameters

*domain-number*: Specifies a PTP domain by its domain number. If the PTP profile is IEEE 1588 version 2 or IEEE 802.1AS, the value range for the *domain-number* argument is 0 to 255. If the PTP profile is SMPTE ST 2059-2, the value range for the *domain-number* argument is 0 to 127.

## Change description

The default value 127 and the value range 0 to 127 were added for specifying the PTP domain number in the SMPTE ST 2059-2 PTP profile.

## Modified command: ptp min-delayreq-interval

### Syntax

```
ptp min-delayreq-interval interval
undo ptp min-delayreq-interval
```

### Views

Layer 2 Ethernet interface view

Layer 3 Ethernet interface view

## Parameters

*interval*: Specifies an exponent, to the power of which base 2 is raised to get the interval in seconds for sending delay request messages. If the PTP profile is IEEE 1588 version 2, the value range for the *interval* argument is -4 to 6. If the PTP profile is SMPTE ST 2059-2, the value range for the *interval* argument is -5 to 4.

## Change description

The value range -5 to 4 was added for specifying the minimum delay request message sending interval in the SMPTE ST 2059-2 PTP profile.

## Modified command: ptp pdelay-req-interval

### Syntax

```
ptp pdelay-req-interval interval
undo ptp pdelay-req-interval
```

### Views

Layer 2 Ethernet interface view

Layer 3 Ethernet interface view

## Parameters

*interval*: Specifies an exponent, to the power of which base 2 is raised to get the interval in seconds for sending peer delay request messages. If the PTP profile is IEEE 1588 version 2, the value range for the *interval* argument is 0 to 5. If the PTP profile is IEEE 802.1AS, the value range for the *interval* argument is -4 to 6. If the PTP profile is SMPTE ST 2059-2, the value range for the *interval* argument is -5 to 4.

## Change description

The value range -5 to 4 was added for specifying the minimum peer delay request message sending interval in the SMPTE ST 2059-2 PTP profile.

## Modified command: ptp syn-interval

### Syntax

```
ptp syn-interval interval
undo ptp syn-interval
```

### Views

Layer 2 Ethernet interface view  
Layer 3 Ethernet interface view

### Parameters

*interval*: Specifies an exponent, to the power of which base 2 is raised to get the interval in seconds for sending Sync messages. If the PTP profile is IEEE 802.1AS, the value range for the *interval* argument is -4 to 6. If the PTP profile is IEEE 1588 version 2, the value range for the *interval* argument is -1 to 1. If the PTP profile is SMPTE ST 2059-2, the value range for the *interval* argument is -5 to -1.

### Change description

The value range -5 to -1 was added for specifying the Sync message sending interval in the SMPTE ST 2059-2 PTP profile.

## Modified feature: Shutting down an interface by using OpenFlow

### Feature change description

From this release, tunnel interfaces, VLAN interfaces, and VSI interfaces do not support the **openflow shutdown** command.

### Command changes

#### Modified command: openflow shutdown

### Syntax

```
openflow shutdown
```

### Old views

Layer 2 Ethernet interface view  
Layer 3 Ethernet interface view  
Layer 2 aggregate interface view  
Layer 3 aggregate Interface view  
Tunnel interface view  
VLAN interface view  
VSI interface view

### New views

Layer 2 Ethernet interface view  
Layer 3 Ethernet interface view

Layer 2 aggregate interface view

Layer 3 aggregate Interface view

### Change description

Before modification: Tunnel interfaces, VLAN interfaces, and VSI interfaces support the **openflow shutdown** command.

After modification: Tunnel interfaces, VLAN interfaces, and VSI interfaces do not support the **openflow shutdown** command.

## Modified feature: Configuring flow sampling of sFlow

### Feature change description

From this release, the sFlow instance was specified for the flow sampling increase.

### Command changes

Modified command: sflow counter collector

#### Old syntax

```
sflow flow collector collector-id
```

```
undo sflow flow collector
```

#### New syntax

```
sflow flow [instance instance-id] collector collector-id
```

```
undo sflow flow [instance instance-id] collector
```

#### Views

Layer 2 Ethernet interface view

Layer 3 Ethernet interface view

#### Parameters

**instance** *instance-id*: sFlow instance ID range is 1 to 4. By default, the instance ID is 1. If this parameter is not configured, sFlow instance 1 is specified for flow sampling.

### Change description

Before modification: Only sFlow Collector can be specified for flow sampling.

After modification: sFlow instance and .sFlow Collector can be specified for flow sampling.

## Modified feature: Configuring counter sampling of sFlow

### Feature change description

From this release, the sFlow instance was specified for the counter sampling increase.

## Command changes

Modified command: sflow counter collector

### Old syntax

```
sflow counter collector collector-id
```

```
undo sflow counter collector
```

### New syntax

```
sflow counter [instance instance-id] collector collector-id
```

```
undo sflow counter [instance instance-id] collector
```

### Views

Layer 2 Ethernet interface view

Layer 3 Ethernet interface view

### Parameters

**instance** *instance-id*: sFlow instance ID range is 1 to 4. By default, the instance ID is 1. If this parameter is not configured, sFlow instance 1 is specified for counter sampling.

### Change description

Before modification: Only sFlow Collector can be specified for counter sampling.

After modification: sFlow instance and .sFlow Collector can be specified for counter sampling.

## Modified feature: Testing the reachability of remote VMs in VXLANs

### Feature change description

Before modification, the system automatically allocates ACL resources to reachability tests after L2VPN is enabled.

After modification, the system does not automatically allocate ACL resources to reachability tests after L2VPN is enabled. To perform reachability tests, you must first enable the reachability test feature by using the **emulate-ping vxlan enable** command to obtain ACL resources.

## Command changes

New command: emulate-ping vxlan enable

Use **emulate-ping vxlan enable** to enable the reachability test feature to allocate ACL resources for testing the reachability of remote VMs in VXLANs.

Use **undo emulate-ping vxlan enable** to disable the reachability test feature to release the ACL resources allocated to reachability tests.

### Syntax

```
emulate-ping vxlan enable
```

```
undo emulate-ping vxlan enable
```

## Default

The reachability test feature is disabled.

## Views

System view

## Predefined user roles

network-admin

## Examples

# Enable the reachability test feature to allocate ACL resources to reachability tests.

```
<Sysname> system-view
```

```
[Sysname] emulate-ping vxlan enable
```

## Related commands

**emulate-ping vxlan**

# Modified feature: Enabling packet statistics for automatically created VXLAN tunnels

## Feature change description

The **destination** *ipv4-address* option was added to the command that enables packet statistics for automatically created VXLAN tunnels.

## Command changes

### Modified command: tunnel statistics vxlan auto

#### Old syntax

```
tunnel statistics vxlan auto
```

```
undo tunnel statistics vxlan auto
```

#### New syntax

```
tunnel statistics vxlan auto [destination ipv4-address]
```

```
undo tunnel statistics vxlan auto [destination ipv4-address]
```

## Views

System view

## Parameters

**destination** *ipv4-address*: Specifies an automatically created VXLAN tunnel by its destination IPv4 address. If you do not specify an automatically created VXLAN tunnel, this command enables packet statistics for all automatically created VXLAN tunnels.

## Change description

Before modification: The command does not support specifying an automatically created VXLAN tunnel by its destination address.

After modification: The command supports using the **destination** *ipv4-address* option to specify an automatically created VXLAN tunnel.

## Modified feature: Disabling flooding for a VSI

### Feature change description

The **all-direction** and **dci** keywords were added to the command that disables flooding for a VSI.

### Command changes

#### Modified command: flooding disable

##### Old syntax

```
flooding disable { all | { broadcast | unknown-multicast | unknown-unicast }
* }
undo flooding disable
```

##### New syntax

```
flooding disable { all | { broadcast | unknown-multicast | unknown-unicast }
* } [all-direction | dci]
undo flooding disable
```

##### Views

VSI view

##### Parameters

**all-direction**: Disables flooding traffic received from an AC or VXLAN tunnel interface to any other ACs and VXLAN tunnel interfaces of the same VSI. If you do not specify this keyword, the command only disables flooding traffic received from ACs to VXLAN tunnel interfaces of the VSI.

**dci**: Disables flooding only to VXLAN-DCI tunnel interfaces. If you do not specify this keyword, the command disables flooding to both VXLAN tunnel interfaces and VXLAN-DCI tunnel interfaces.

##### Change description

Before modification: The **all-direction** and **dci** keywords were not supported by this command.

After modification: The **all-direction** and **dci** keywords were added to this command.

## Modified feature: Setting the maximum bandwidth for an AC

### Feature change description

The value range for the maximum bandwidth was modified in an AC.

## Command changes

Modified command: `bandwidth`

### Syntax

```
bandwidth bandwidth
undo bandwidth
```

### Views

Ethernet service instance view

### Change description

Before modification: The value range for the *bandwidth* argument is 64 to 4194303 kbps.

After modification: The value range for the *bandwidth* argument is 64 to 167772159 kbps.

## Modified feature: Displaying information about peers automatically discovered through BGP in EVPN

### Feature change description

Support for displaying information about peers discovered through MAC/IP advertisement routes was added.

## Command changes

Modified command: `display evpn auto-discovery`

### Old syntax

```
display evpn auto-discovery { imet [peer ip-address] [vsi vsi-name] |
macip-prefix [nexthop next-hop] [count] }
```

### New syntax

```
display evpn auto-discovery { imet [peer ip-address] [vsi vsi-name] |
mac-ip | macip-prefix [nexthop next-hop] [count] }
```

### Views

System view

### Parameters

**mac-ip**: Specifies peers discovered through MAC/IP advertisement routes.

### Examples

# Display information about peers discovered through MAC/IP advertisement routes.

```
<Sysname> display evpn auto-discovery mac-ip
```

Total number of automatically discovered peers: 1

VSI name: vpna

| Destination IP | Source IP | VXLAN ID | Tunnel mode | Tunnel name |
|----------------|-----------|----------|-------------|-------------|
| 6.6.6.6        | 1.1.1.9   | 100      | VXLAN       | Tunnel1     |

**Table 68 Command output**

| Field          | Description                      |
|----------------|----------------------------------|
| Destination IP | Tunnel destination IP address.   |
| Source IP      | Tunnel source IP address.        |
| Tunnel mode    | Tunnel mode, which can be VXLAN. |

### Change description

Before modification: The **mac-ip** keyword was not supported by this command.

After modification: The **mac-ip** keyword was added to this command.

## Modified feature: PW redundancy

### Feature change description

Support for the master/slave PW redundancy mode was added.

In master/slave mode, one endpoint PE of a PW operates as the master node and the other endpoint PE operates as the slave node. The master PE determines the active/standby state of the PW and uses LDP to advertise the PW state to the slave PE. The slave PE uses the same PW state as the master PE based on the information received from the master PE. In this way, the master and slave PEs for the set of redundant PWs can use the same active PW to forward user traffic.

The slave PE does not need to send the PW states to the master PE and the master PE ignores the PW states sent by the slave PE.

### Command changes

#### New command: pw-redundancy

Use **pw-redundancy** to specify the master/slave PW redundancy mode and configure the local PE as the master node.

Use **undo pw-redundancy** to restore the default.

#### Syntax

**pw-redundancy master**

**undo pw-redundancy**

#### Default

The PW redundancy mode is master/slave and the local PE operates as the slave node.

#### Views

Cross-connect view

VSI LDP signaling view

#### Predefined user roles

network-admin

#### Parameters

**master**: Uses the master/slave PW redundancy mode and configures the local PE as the master node.

## Usage guidelines

This command is applicable only to LDP PWs.

For LDP PWs, PEs use LDP to negotiate the active/standby states of the PWs. In master/slave mode, one endpoint PE of a PW operates as the master node and the other endpoint PE operates as the slave node. The master PE determines the active/standby state of the PW and uses LDP to advertise the PW state to the slave PE. The slave PE uses the same PW state as the master PE based on the information received from the master PE. In this way, the master and slave PEs for the set of redundant PWs can use the same active PW to forward user traffic.

The slave PE does not need to send the PW states to the master PE and the master PE ignores the PW states sent by the slave PE.

Do not configure this command on the local PE if the remote PE does not support the master/slave PW redundancy mode.

## Examples

```
Use the master/slave PW redundancy mode and configure the local PE as the master node.
```

```
<Sysname> system-view
[Sysname] xconnect-group vpn1
[Sysname-xcg-vpn1] connection pw2pw
[Sysname-xcg-vpn1-pw2pw] pw-redundancy master
```

## Modified command: display l2vpn ldp

### Syntax

```
display l2vpn ldp [peer ip-address [pw-id pw-id] | xconnect-group group-name] [verbose]
```

### Views

Any view

### Change description

The **PW Preferential Forwarding Status Bit** field was added to the **display l2vpn ldp verbose** command output.

This field indicates the PW state processing mode. Possible values are:

- **Process**—Process the received Active/Standby status bits.
- **Ignore**—Ignore the received Active/Standby status bits.

For example:

```
Display detailed information about all PW labels advertised by LDP.
```

```
<Sysname> display l2vpn ldp verbose
Peer: 192.2.2.2 PW ID: 1000
 Xconnect-group: vpn1
 Connection : ldp
 PW State : Up
 PW Status Communication: Notification method
 PW Preferential Forwarding Status Bit: Process
 PW ID FEC (Local/Remote):
 PW Type : VLAN/VLAN
 Group ID : 0/0
 Label : 1151/1279
 Control Word : Disabled/Disabled
```

```
VCCV CC Type: -/-
VCCV CV Type: -/-
MTU : 1500/1500
PW Status : PW forwarding/PW forwarding
```

## Modified command: display l2vpn xconnect-group

### Syntax

```
display l2vpn xconnect-group [name group-name] [verbose]
```

### Views

Any view

### Change description

The **PW Redundancy** field was added to the **display l2vpn xconnect-group verbose** command output.

This field indicates the PW redundancy mode. Possible values are:

- **Slave**—Master/slave mode and the local PE operates as the slave node.
- **Master**—Master/slave mode and the local PE operates as the master node.

For example:

# Display detailed information about all cross-connect groups.

```
<Sysname> display l2vpn xconnect-group verbose
```

```
Xconnect-group Name: vpnb
```

```
Connection of auto-discovery: Site 1, Remote Site 2
```

```
Connection ID : 0
```

```
State : Up
```

```
MTU : 1500
```

```
PW Redundancy : Slave
```

```
BGP PWs:
```

| Peer      | Remote Site | Link ID | State |
|-----------|-------------|---------|-------|
| 192.3.3.3 | 2           | 1       | Up    |

```
ACs:
```

| AC            | Link ID | State | Type   |
|---------------|---------|-------|--------|
| XGE1/0/1 srv1 | 0       | Up    | Manual |

## Modified command: display l2vpn vsi

### Syntax

```
display l2vpn vsi [name vsi-name] [verbose]
```

### Views

Any view

### Change description

The **PW Redundancy** field was added to the **display l2vpn vsi verbose** command output.

This field indicates the PW redundancy mode. Possible values are:

- **Slave**—Master/slave mode and the local PE operates as the slave node.
- **Master**—Master/slave mode and the local PE operates as the master node.

For example:

# Display detailed information about all VSIs.

```
<Sysname> display l2vpn vsi verbose
```

VSI Name: vpls1

```
VSI Index : 0
VSI Description : vsi for vpls1
VSI State : Up
MTU : 1500
Bandwidth : Unlimited
Broadcast Restrain : Unlimited
Multicast Restrain : Unlimited
Unknown Unicast Restrain: Unlimited
MAC Learning : -
MAC Table Limit : -
Drop Unknown : -
PW Redundancy : Master
Flooding : Enabled
Statistics : Disabled
VXLAN ID : -
```

LDP PWs:

| Peer      | PW ID | Link ID | State   |
|-----------|-------|---------|---------|
| 192.3.3.3 | 1     | 8       | Up      |
| 192.3.3.3 | 1001  | 8       | Blocked |

BGP PWs:

| Peer      | Remote Site | Link ID | State |
|-----------|-------------|---------|-------|
| 192.4.4.4 | 1           | 9       | Up    |

ACs:

| AC            | Link ID | State | Type   |
|---------------|---------|-------|--------|
| Vlan10        | 0       | Up    | Manual |
| XGE1/0/1 srv1 | 1       | Up    | Manual |

## Modified command: peer

### Old syntax

In cross-connect view:

```
peer ip-address pw-id pw-id [in-label label-value out-label label-value]
[pw-class class-name | tunnel-policy tunnel-policy-name] *
```

```
undo peer ip-address pw-id pw-id
```

In VSI LDP signaling view:

```
peer ip-address [pw-id pw-id] [no-split-horizon | pw-class class-name |
tunnel-policy tunnel-policy-name] *
```

```
undo peer ip-address pw-id pw-id
```

### New syntax

In cross-connect view:

```
peer ip-address pw-id pw-id [ignore-standby-state | in-label label-value
out-label label-value] [pw-class class-name | tunnel-policy
tunnel-policy-name] *
```

```
undo peer ip-address pw-id pw-id
```

In VSI LDP signaling view:

```
peer ip-address [pw-id pw-id] [ignore-standby-state | no-split-horizon
| pw-class class-name | tunnel-policy tunnel-policy-name] *
```

```
undo peer ip-address pw-id pw-id
```

## Views

Cross-connect view

VSI LDP signaling view

## Change description

The **ignore-standby-state** keyword was added to the peer command syntax.

This keyword enables the device to ignore the PW active/standby status bit received from the remote PE. If you do not specify this keyword, the device does not ignore the PW active/standby status bit received from the remote PE.

This keyword is applicable only to LDP PWs.

# Modified feature: Displaying log buffer information and buffered logs

## Feature change description

The **last-mins mins** option was added to the **display logbuffer** command. You can use the option to display log buffer information and logs buffered over the last specified period of time.

## Command changes

### Modified command: display logbuffer

#### Old syntax

```
display logbuffer [module module-name] [reverse] [level severity | size
buffer-size | slot slot-number] *
```

#### New syntax

```
display logbuffer [module module-name] [reverse] [level severity | size
buffer-size | slot slot-number] * [last-mins mins]
```

## Views

Any view

## Change description

The **last-mins mins** option was added to the command. This option specifies a time period in minutes. The value range is 1 to 43200.

If you specify a time period by using this option, the command displays logs buffered over the last specified time period.

If you do not specify a time period, the command displays all logs in the log buffer.

# Modified feature: Configuring the timestamp for logs sent to log hosts

## Feature change description

The **with-milliseconds** keyword was added to the **info-center timestamp loghost** command. You can set the timestamp to be accurate to milliseconds for logs output to log hosts in date or ISO 8601 format.

## Command changes

### Modified command: display logbuffer

#### Old syntax

```
info-center timestamp loghost { date | iso [with-timezone] | no-year-date
| none }
```

#### New syntax

```
info-center timestamp loghost { date [with-milliseconds] | iso
[with-milliseconds | with-timezone] * | no-year-date | none }
```

#### Views

System view

#### Change description

The **with-milliseconds** keyword was added to the command.

This keyword sets the timestamp to be accurate to milliseconds. The millisecond value is appended to the time information in the timestamp with a dot as the separator.

- Example of a timestamp in date format with millisecond accuracy: Dec 8 10:12:21.708 2018.
- Example of a timestamp in ISO 8601 format with millisecond accuracy:  
2018-09-21T15:32:55.708.

If you do not specify this keyword, the timestamp in date or ISO 8601 format is accurate to seconds.

# Modified feature: Configuring log output to a log host

## Feature change description

The **filter** *filter-name* option was added to the **info-center loghost** command. You can use this option to specify a log output filter to filter the logs to be output to the log host.

## Command changes

### Modified command: info-center loghost

#### Old syntax

```
info-center loghost [vpn-instance vpn-instance-name] { hostname |
ipv4-address | ipv6 ipv6-address } [port port-number] [dscp dscp-value]
[facility local-number]
```

## New syntax

```
info-center loghost [vpn-instance vpn-instance-name] { hostname |
ipv4-address | ipv6 ipv6-address } [port port-number] [dscp dscp-value]
[facility local-number] [filter filter-name]
```

## Views

System view

## Change description

The **filter** *filter-name* option was added to the command. This option specifies a log output filter by the filter name, a case-insensitive string of 1 to 8 characters.

If you specify a log output filter by using this option, the system controls log output to the log host according to the rules set in the specified log output filter.

If you do not specify a log output filter, the log output rules configured by using the **info-center source** command for the log host destination are used.

# Modified feature: IRF master election during automated VCF fabric deployment

## Feature change description

From this software version, two IRF member devices in an IRF fabric use the following rules to elect the IRF master during automated VCF fabric deployment:

- If the uptime of both devices is shorter than two hours, the device with the higher bridge MAC address becomes the IRF master.
- If the uptime of one device is equal to or longer than two hours, that device becomes the IRF master.
- If the uptime of both devices are longer than two hours, the IRF fabric cannot be set up. You must manually reboot one of the member devices. The rebooted device will become the IRF subordinate.

## Command changes

None.