



**7450 ETHERNET SERVICE SWITCH
7750 SERVICE ROUTER
7950 EXTENSIBLE ROUTING SYSTEM
VIRTUALIZED SERVICE ROUTER**

**ROUTER CONFIGURATION GUIDE
RELEASE 15.0.R5**

3HE 11976 AAAC TQZZA 01

Issue: 01

September 2017

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

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

© 2017 Nokia.

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

Table of Contents

| | | |
|----------|--|-----------|
| 1 | Getting Started | 11 |
| 1.1 | About This Guide | 11 |
| 1.2 | Router Configuration Process | 13 |
| | | |
| 2 | IP Router Configuration | 15 |
| 2.1 | Configuring IP Router Parameters | 15 |
| 2.1.1 | Interfaces | 15 |
| 2.1.1.1 | Network Interface | 15 |
| 2.1.1.2 | Network Domains | 16 |
| 2.1.1.3 | System Interface | 17 |
| 2.1.1.4 | Unicast Reverse Path Forwarding Check (uRPF) | 17 |
| 2.1.1.5 | Creating an IP Address Range | 18 |
| 2.1.1.6 | QoS Policy Propagation Using BGP (QPPB) | 19 |
| 2.1.1.7 | QPPB | 21 |
| 2.1.1.8 | QPPB and GRT Lookup | 26 |
| 2.1.2 | Router ID | 28 |
| 2.1.3 | Autonomous Systems (AS) | 29 |
| 2.1.4 | Confederations | 30 |
| 2.1.5 | Proxy ARP | 31 |
| 2.1.6 | Exporting an Inactive BGP Route from a VPRN | 32 |
| 2.1.7 | DHCP Relay | 32 |
| 2.1.8 | Internet Protocol Versions | 32 |
| 2.1.8.1 | IPv6 Address Format | 34 |
| 2.1.8.2 | IPv6 Applications | 35 |
| 2.1.8.3 | DNS | 37 |
| 2.1.8.4 | Secure Neighbor Discovery (SeND) | 37 |
| 2.1.8.5 | SeND Persistent CGAs | 39 |
| 2.1.8.6 | IPv6 Provider Edge Router over MPLS (6PE) | 44 |
| 2.1.9 | Static Route Resolution Using Tunnels | 46 |
| 2.1.9.1 | Static Route ECMP Support | 48 |
| 2.2 | Weighted Load Balancing over MPLS LSP | 49 |
| 2.2.1 | Weighted Load Balancing IGP, BGP, and Static Route Prefix Packets over IGP Shortcut | 50 |
| 2.2.1.1 | Feature Configuration | 50 |
| 2.2.1.2 | Feature Behavior | 51 |
| 2.2.1.3 | ECMP Considerations | 52 |
| 2.2.1.4 | Weighted Load Balancing Static Route Packets over MPLS LSP | 52 |
| 2.2.2 | Weighted Load Balancing for 6PE | 54 |
| 2.3 | Class-Based Forwarding of IPv4/IPv6 Prefix Over IGP IPv4 Shortcut | 55 |
| 2.3.1 | Feature Configuration | 55 |
| 2.3.2 | Feature Behavior | 56 |
| 2.3.3 | Feature Limitations | 58 |
| 2.3.4 | Data Path Support | 59 |
| 2.3.5 | Example Configuration and Default CBF Set Election | 60 |

| | | |
|----------|---|----|
| 2.4 | Bidirectional Forwarding Detection..... | 64 |
| 2.4.1 | BFD Control Packet..... | 64 |
| 2.4.2 | Control Packet Format..... | 65 |
| 2.4.3 | BFD for RSVP-TE..... | 66 |
| 2.4.4 | Echo Support..... | 67 |
| 2.4.5 | BFD Support for BGP..... | 67 |
| 2.4.6 | Centralized BFD..... | 67 |
| 2.4.6.1 | IES Over Spoke SDP..... | 67 |
| 2.4.6.2 | BFD Over LAG and VSM Interfaces..... | 68 |
| 2.4.6.3 | LSP BFD and VCCV BFD..... | 69 |
| 2.4.7 | Aggregate Next Hop..... | 69 |
| 2.4.8 | Invalidate Next-Hop Based on ARP/Neighbor Cache State..... | 70 |
| 2.4.8.1 | Invalidate Next-Hop Based on IPV4 ARP..... | 70 |
| 2.4.8.2 | Invalidate Next-Hop Based on Neighbor Cache State..... | 70 |
| 2.4.9 | LDP Shortcut for IGP Route Resolution..... | 71 |
| 2.4.9.1 | IGP Route Resolution..... | 71 |
| 2.4.9.2 | LDP-IGP Synchronization..... | 72 |
| 2.4.9.3 | LDP Shortcut Forwarding Plane..... | 72 |
| 2.4.9.4 | ECMP Considerations..... | 73 |
| 2.4.9.5 | Handling of Control Packets..... | 73 |
| 2.4.9.6 | Handling of Multicast Packets..... | 74 |
| 2.4.9.7 | Interaction with BGP Route Resolution to an LDP FEC..... | 74 |
| 2.4.9.8 | Interaction with Static Route Resolution to an LDP FEC..... | 74 |
| 2.4.9.9 | LDP Control Plane..... | 74 |
| 2.5 | Weighted Load-Balancing over Interface Next-hops..... | 76 |
| 2.6 | GRE Tunnel Overview..... | 77 |
| 2.6.1 | Sample GRE Tunnel Configurations..... | 78 |
| 2.7 | Router Interface Encryption with NGE..... | 79 |
| 2.7.1 | NGE Domains..... | 80 |
| 2.7.1.1 | Private IP/MPLS Network NGE Domain..... | 81 |
| 2.7.1.2 | Private Over Intermediary Network NGE Domain..... | 82 |
| 2.7.2 | Router Interface NGE Domain Concepts..... | 84 |
| 2.7.3 | GRE-MPLS and MPLSoUDP Packets Inside the NGE Domain..... | 85 |
| 2.7.4 | EVPN-VXLAN Tunnels and Services..... | 86 |
| 2.7.5 | Router Encryption Exceptions using ACLs..... | 86 |
| 2.7.6 | IPSec Packets Crossing an NGE Domain..... | 88 |
| 2.7.7 | Multicast Packets Traversing the NGE Domain..... | 89 |
| 2.7.8 | Assigning Key Groups to Router Interfaces..... | 91 |
| 2.7.9 | NGE and BFD Support..... | 91 |
| 2.7.10 | NGE and ACL Interactions..... | 91 |
| 2.7.11 | Router Interface NGE and ICMP Interactions Over the NGE Domain..... | 92 |
| 2.7.12 | 1588v2 Encryption With NGE..... | 93 |
| 2.8 | Process Overview..... | 94 |
| 2.9 | Configuration Notes..... | 95 |
| 2.10 | Configuring an IP Router with CLI..... | 97 |
| 2.10.1 | Router Configuration Overview..... | 97 |
| 2.10.1.1 | System Interface..... | 97 |
| 2.10.1.2 | Network Interface..... | 98 |

| | | |
|----------|---|------------|
| 2.10.2 | Basic Configuration | 98 |
| 2.10.3 | Common Configuration Tasks | 99 |
| 2.10.3.1 | Configuring a System Name..... | 99 |
| 2.10.3.2 | Configuring Interfaces | 99 |
| 2.10.3.3 | Deriving the Router ID | 115 |
| 2.10.3.4 | Configuring a Confederation..... | 116 |
| 2.10.3.5 | Configuring an Autonomous System..... | 117 |
| 2.10.3.6 | Configuring Overload State on a Single SFM..... | 117 |
| 2.11 | Service Management Tasks | 119 |
| 2.11.1 | Changing the System Name..... | 119 |
| 2.11.2 | Modifying Interface Parameters..... | 119 |
| 2.11.3 | Removing a Key Group from a Router Interface | 120 |
| 2.11.4 | Changing the Key Group for a Router Interface | 121 |
| 2.11.5 | Deleting a Logical IP Interface..... | 122 |
| 2.12 | IP Router Configuration Command Reference | 123 |
| 2.12.1 | Command Hierarchies | 123 |
| 2.12.1.1 | Router Commands | 123 |
| 2.12.1.2 | Router BFD Commands | 126 |
| 2.12.1.3 | Router L2TP Commands..... | 126 |
| 2.12.1.4 | Router Interface Commands | 130 |
| 2.12.1.5 | Router Interface IPv6 Commands | 133 |
| 2.12.1.6 | Router Advertisement Commands | 134 |
| 2.12.2 | Command Descriptions | 135 |
| 2.12.2.1 | Generic Commands..... | 135 |
| 2.12.2.2 | Router Global Commands | 137 |
| 2.12.2.3 | Router L2TP Commands..... | 181 |
| 2.12.2.4 | Router Interface Commands | 212 |
| 2.12.2.5 | Router Interface Encryption Commands | 277 |
| 2.12.2.6 | Router Advertisement Commands | 279 |
| 2.13 | Show, Clear, and Debug Command Reference | 287 |
| 2.13.1 | Command Hierarchies..... | 287 |
| 2.13.1.1 | Show Commands | 287 |
| 2.13.1.2 | Clear Commands..... | 289 |
| 2.13.1.3 | Debug Commands..... | 290 |
| 2.13.1.4 | Tools Commands | 290 |
| 2.13.2 | Command Descriptions | 291 |
| 2.13.2.1 | Show Commands | 291 |
| 2.13.2.2 | Clear Commands..... | 420 |
| 2.13.2.3 | Debug Commands..... | 427 |
| 2.13.2.4 | Tools Commands | 432 |
| 3 | VRRP | 441 |
| 3.1 | VRRP Overview..... | 441 |
| 3.2 | VRRP Components | 442 |
| 3.2.1 | Virtual Router..... | 442 |
| 3.2.2 | IP Address Owner | 442 |
| 3.2.3 | Primary and Secondary IP Addresses..... | 443 |
| 3.2.4 | Virtual Router Master..... | 443 |
| 3.2.5 | Virtual Router Backup..... | 444 |

| | | |
|----------|---|-----|
| 3.2.6 | Owner and Non-Owner VRRP..... | 444 |
| 3.2.7 | Configurable Parameters..... | 444 |
| 3.2.7.1 | Virtual Router ID (VRID)..... | 445 |
| 3.2.7.2 | Priority | 445 |
| 3.2.7.3 | IP Addresses | 446 |
| 3.2.7.4 | Message Interval and Master Inheritance | 446 |
| 3.2.7.5 | Skew Time..... | 447 |
| 3.2.7.6 | Master Down Interval..... | 447 |
| 3.2.7.7 | Preempt Mode..... | 448 |
| 3.2.7.8 | VRRP Message Authentication | 448 |
| 3.2.7.9 | Authentication Data | 450 |
| 3.2.7.10 | Virtual MAC Address | 451 |
| 3.2.7.11 | VRRP Advertisement Message IP Address List Verification | 451 |
| 3.2.7.12 | Inherit Master VRRP Router's Advertisement Interval Timer | 452 |
| 3.2.7.13 | IPv6 Virtual Router Instance Operationally Up | 452 |
| 3.2.7.14 | Policies | 452 |
| 3.3 | VRRP Priority Control Policies | 453 |
| 3.3.1 | VRRP Virtual Router Policy Constraints..... | 453 |
| 3.3.2 | VRRP Virtual Router Instance Base Priority..... | 453 |
| 3.3.3 | VRRP Priority Control Policy Delta In-Use Priority Limit | 454 |
| 3.3.4 | VRRP Priority Control Policy Priority Events | 454 |
| 3.3.4.1 | Priority Event Hold-Set Timers | 455 |
| 3.3.4.2 | Port Down Priority Event | 455 |
| 3.3.4.3 | LAG Degrade Priority Event | 456 |
| 3.3.4.4 | Host Unreachable Priority Event | 458 |
| 3.3.4.5 | Route Unknown Priority Event..... | 458 |
| 3.4 | VRRP Non-Owner Accessibility..... | 460 |
| 3.4.1 | Non-Owner Access Ping Reply | 460 |
| 3.4.2 | Non-Owner Access Telnet..... | 460 |
| 3.4.3 | Non-Owner Access SSH | 461 |
| 3.5 | VRRP Configuration Process Overview | 462 |
| 3.6 | Configuration Notes..... | 464 |
| 3.6.1 | General..... | 464 |
| 3.7 | Configuring VRRP with CLI | 465 |
| 3.7.1 | VRRP Configuration Overview | 465 |
| 3.7.1.1 | Preconfiguration Requirements | 465 |
| 3.7.2 | Basic VRRP Configurations..... | 465 |
| 3.7.2.1 | VRRP Policy | 466 |
| 3.7.2.2 | VRRP IES Service Parameters | 467 |
| 3.7.2.3 | VRRP Router Interface Parameters | 469 |
| 3.7.3 | Common Configuration Tasks | 469 |
| 3.7.3.1 | Creating Interface Parameters | 470 |
| 3.7.4 | Configuring VRRP Policy Components | 471 |
| 3.7.4.1 | Configuring Service VRRP Parameters..... | 471 |
| 3.7.4.2 | Configuring Router Interface VRRP Parameters..... | 472 |
| 3.8 | VRRP Configuration Management Tasks..... | 474 |
| 3.8.1 | Modifying a VRRP Policy..... | 474 |
| 3.8.1.1 | Deleting a VRRP Policy..... | 474 |
| 3.8.2 | Modifying Service and Interface VRRP Parameters..... | 475 |

| | | |
|----------|--|------------|
| 3.8.2.1 | Modifying Non-Owner Parameters | 475 |
| 3.8.2.2 | Modifying Owner Parameters | 475 |
| 3.8.2.3 | Deleting VRRP from an Interface or Service | 475 |
| 3.9 | VRRP Configuration Command Reference | 477 |
| 3.9.1 | Command Hierarchies | 477 |
| 3.9.1.1 | IPv4 Interface VRRP Commands | 477 |
| 3.9.1.2 | Router Interface Commands | 478 |
| 3.9.1.3 | IPv6 Interface VRRP Commands | 478 |
| 3.9.1.4 | Priority Control Event Policy Commands | 479 |
| 3.9.2 | Command Descriptions | 480 |
| 3.9.2.1 | Interface Configuration Commands | 481 |
| 3.9.2.2 | Priority Policy Commands | 501 |
| 3.9.2.3 | Priority Policy Event Commands | 503 |
| 3.9.2.4 | Priority Policy Port Down Event Commands | 507 |
| 3.9.2.5 | Priority Policy LAG Events Commands | 510 |
| 3.9.2.6 | Priority Policy Host Unreachable Event Commands | 512 |
| 3.9.2.7 | Priority Policy Route Unknown Event Commands | 517 |
| 3.10 | Show, Monitor, Clear, and Debug Command Reference | 525 |
| 3.10.1 | Command Hierarchies | 525 |
| 3.10.1.1 | Show Commands | 525 |
| 3.10.1.2 | Monitor Commands | 525 |
| 3.10.1.3 | Clear Commands | 526 |
| 3.10.1.4 | Debug Commands | 526 |
| 3.10.2 | Command Descriptions | 526 |
| 3.10.2.1 | Show Commands | 526 |
| 3.10.2.2 | Monitor Commands | 541 |
| 3.10.2.3 | Clear Commands | 543 |
| 3.10.2.4 | Debug Commands | 544 |
| 4 | Filter Policies | 547 |
| 4.1 | ACL Filter Policy Overview | 547 |
| 4.1.1 | Filter Policy Basics | 548 |
| 4.1.1.1 | Filter Policy Packet Match Criteria | 549 |
| 4.1.1.2 | IPv4/IPv6 Filter Policy Entry Match Criteria | 549 |
| 4.1.1.3 | MAC Filter Policy Entry Match Criteria | 551 |
| 4.1.1.4 | IP Exception Filters | 552 |
| 4.1.1.5 | Filter Policy Actions | 553 |
| 4.1.1.6 | Viewing Filter Policy Actions | 559 |
| 4.1.1.7 | Filter Policy Statistics | 561 |
| 4.1.1.8 | Filter Policy Logging | 562 |
| 4.1.1.9 | Filter Policy cflowd Sampling | 563 |
| 4.1.1.10 | Filter Policy Management | 563 |
| 4.1.2 | Filter Policy Advanced Topics | 564 |
| 4.1.2.1 | Match List for Filter Policies | 564 |
| 4.1.2.2 | Embedded Filters | 567 |
| 4.1.2.3 | System-level IPv4/IPv6 Line Card Filter Policy | 569 |
| 4.1.2.4 | Primary and Secondary Filter Policy Action for PBR/PBF Redundancy | 570 |
| 4.1.2.5 | Extended Action for Performing Two Actions at a Time | 572 |

| | | |
|----------|---|-----|
| 4.1.2.6 | Advanced VPRN Redirection | 573 |
| 4.1.2.7 | Destination MAC Rewrite When Deploying Policy-Based Forwarding..... | 574 |
| 4.1.2.8 | Network-port VPRN Filter Policy | 576 |
| 4.1.2.9 | ISID MAC Filters | 576 |
| 4.1.2.10 | VID MAC Filters..... | 577 |
| 4.1.2.11 | IP Exception Filters..... | 580 |
| 4.1.2.12 | Redirect Policies..... | 581 |
| 4.1.2.13 | HTTP-redirect (Captive Portal)..... | 583 |
| 4.1.2.14 | Filter Policies and Dynamic Policy-Driven Interfaces | 585 |
| 4.1.2.15 | Filter Policy-based ESM Service Chaining | 586 |
| 4.1.2.16 | Policy-Based Forwarding for Deep Packet Inspection in VPLS | 591 |
| 4.2 | Configuring Filter Policies with CLI..... | 595 |
| 4.2.1 | Common Configuration Tasks | 595 |
| 4.2.1.1 | Creating an IPv4 Filter Policy | 595 |
| 4.2.1.2 | Creating an IPv6 Filter Policy | 598 |
| 4.2.1.3 | Creating a MAC Filter Policy | 598 |
| 4.2.1.4 | Creating an IP Exception Filter Policy | 600 |
| 4.2.1.5 | Creating a Match List for Filter Policies | 602 |
| 4.2.1.6 | Applying Filter Policies | 602 |
| 4.2.1.7 | Creating a Redirect Policy | 605 |
| 4.3 | Filter Management Tasks..... | 607 |
| 4.3.1 | Renumbering Filter Policy Entries | 607 |
| 4.3.2 | Modifying a Filter Policy..... | 609 |
| 4.3.3 | Deleting a Filter Policy..... | 610 |
| 4.3.4 | Modifying a Redirect Policy | 610 |
| 4.3.5 | Deleting a Redirect Policy | 612 |
| 4.3.6 | Copying Filter Policies | 612 |
| 4.4 | Filter Configuration Command Reference | 615 |
| 4.4.1 | Command Hierarchies..... | 615 |
| 4.4.1.1 | IPv4 Filter Policy Commands | 615 |
| 4.4.1.2 | IPv6 Filter Policy Commands | 618 |
| 4.4.1.3 | MAC Filter Commands | 621 |
| 4.4.1.4 | IP Exception Filter Policy Configuration Commands | 622 |
| 4.4.1.5 | System Filter Policy Commands..... | 622 |
| 4.4.1.6 | Redirect Policy Configuration Commands..... | 623 |
| 4.4.1.7 | Match Filter List Commands..... | 624 |
| 4.4.1.8 | Log Filter Commands | 624 |
| 4.4.1.9 | Copy Filter Commands..... | 625 |
| 4.4.2 | Command Descriptions | 625 |
| 4.4.2.1 | Generic Commands..... | 626 |
| 4.4.2.2 | Global Filter Commands..... | 626 |
| 4.4.2.3 | Filter Log Commands | 629 |
| 4.4.2.4 | ACL Filter Policy Commands..... | 631 |
| 4.4.2.5 | General Filter Entry Commands | 640 |
| 4.4.2.6 | IP (v4/v6) and IP Exception Filter Entry Commands | 643 |
| 4.4.2.7 | Match List Configuration Commands | 670 |
| 4.4.2.8 | MAC Filter Entry Commands | 675 |
| 4.4.2.9 | MAC Filter Match Criteria | 676 |

| | | |
|----------|---|------------|
| 4.4.2.10 | IP Exception Filter Policy Commands | 684 |
| 4.4.2.11 | Policy and Entry Maintenance Commands..... | 685 |
| 4.4.2.12 | Redirect Policy Commands | 686 |
| 4.5 | Show, Clear, Monitor, and Debug Command Reference | 695 |
| 4.5.1 | Command Hierarchies | 695 |
| 4.5.1.1 | Show Commands | 695 |
| 4.5.1.2 | Clear Commands..... | 696 |
| 4.5.1.3 | Monitor Commands | 696 |
| 4.5.1.4 | Debug Commands..... | 696 |
| 4.5.1.5 | Tools Commands | 697 |
| 4.5.2 | Command Descriptions | 697 |
| 4.5.2.1 | Show Commands | 697 |
| 4.5.2.2 | Clear Commands..... | 745 |
| 4.5.2.3 | Monitor Commands | 747 |
| 4.5.2.4 | Debug Commands..... | 750 |
| 5 | Hybrid OpenFlow Switch | 759 |
| 5.1 | In This Chapter | 759 |
| 5.2 | Hybrid OpenFlow Switching | 760 |
| 5.2.1 | Redundant Controllers and Multiple Switch Instances | 762 |
| 5.2.2 | GRT-only and Multi-Service H-OFS Modes of Operations..... | 763 |
| 5.2.2.1 | Port and VLAN ID Match in Flow Table Entries..... | 765 |
| 5.2.3 | Hybrid OpenFlow Switch Steering using Filter Policies | 766 |
| 5.2.4 | Hybrid OpenFlow Switch Statistics..... | 769 |
| 5.2.5 | OpenFlow Switch Auxiliary Channels..... | 770 |
| 5.2.6 | Hybrid OpenFlow Switch Traffic Steering Details..... | 771 |
| 5.2.6.1 | SR OS H-OFS Logical Port | 771 |
| 5.2.6.2 | SR OS H-OFS Port and VLAN Encoding | 772 |
| 5.2.6.3 | Redirect to IP next-hop..... | 775 |
| 5.2.6.4 | Redirect to GRT Instance or VRF Instance | 776 |
| 5.2.6.5 | Redirect to Next-hop and VRF/GRT Instance | 776 |
| 5.2.6.6 | Redirect to ESI (L2)..... | 777 |
| 5.2.6.7 | Redirect to ESI (L3)..... | 777 |
| 5.2.6.8 | Redirect to ESI IP VAS-Interface Router..... | 778 |
| 5.2.6.9 | Redirect to LSP | 779 |
| 5.2.6.10 | Redirect to NAT | 779 |
| 5.2.6.11 | Redirect to SAP | 780 |
| 5.2.6.12 | Redirect to SDP..... | 780 |
| 5.2.6.13 | Redirect to a Specific LSP Used by a VPRN Service..... | 781 |
| 5.2.6.14 | Forward Action | 782 |
| 5.2.6.15 | Drop Action..... | 783 |
| 5.2.6.16 | Default No-match Action..... | 783 |
| 5.2.6.17 | Programming of DSCP Remark Action | 783 |
| 5.3 | Configuration Notes..... | 785 |
| 5.4 | OpenFlow Command Reference | 787 |
| 5.4.1 | Command Hierarchies..... | 787 |
| 5.4.1.1 | OpenFlow Commands..... | 787 |
| 5.4.1.2 | Show Commands | 787 |
| 5.4.1.3 | Tools Commands | 788 |

| | | |
|----------|--|------------|
| 5.4.2 | Command Descriptions | 788 |
| 5.4.2.1 | Generic Commands..... | 788 |
| 5.4.2.2 | Show Commands | 793 |
| 5.4.2.3 | Debug Commands..... | 799 |
| 6 | Cflowd | 811 |
| 6.1 | Cflowd Overview..... | 811 |
| 6.1.1 | Operation..... | 811 |
| 6.1.1.1 | Version 8 | 814 |
| 6.1.1.2 | Version 9 | 814 |
| 6.1.1.3 | Version 10 | 814 |
| 6.1.2 | Cflowd Filter Matching..... | 815 |
| 6.2 | Cflowd Configuration Process Overview | 816 |
| 6.3 | Configuration Notes..... | 817 |
| 6.4 | Configuring Cflowd with CLI | 819 |
| 6.4.1 | Cflowd Configuration Overview | 819 |
| 6.4.1.1 | Traffic Sampling..... | 819 |
| 6.4.1.2 | Collectors..... | 820 |
| 6.4.2 | Basic Cflowd Configuration | 821 |
| 6.4.3 | Common Configuration Tasks | 822 |
| 6.4.3.1 | Global Cflowd Components..... | 822 |
| 6.4.3.2 | Enabling Cflowd..... | 823 |
| 6.4.3.3 | Configuring Global Cflowd Parameters | 823 |
| 6.4.3.4 | Configuring Cflowd Collectors | 824 |
| 6.4.3.5 | Specifying Cflowd Options on an IP Interface | 836 |
| 6.4.3.6 | Specifying Sampling Options in Filter Entries..... | 837 |
| 6.5 | Cflowd Configuration Management Tasks..... | 840 |
| 6.5.1 | Modifying Global Cflowd Components | 840 |
| 6.5.2 | Modifying Cflowd Collector Parameters | 841 |
| 6.6 | Cflowd Configuration Command Reference | 843 |
| 6.6.1 | Command Hierarchies..... | 843 |
| 6.6.2 | Command Descriptions | 844 |
| 6.6.2.1 | Global Commands..... | 844 |
| 6.7 | Show, Tools, and Clear Command Reference | 857 |
| 6.7.1 | Command Hierarchies..... | 857 |
| 6.7.1.1 | Show Commands | 857 |
| 6.7.1.2 | Tools Commands | 857 |
| 6.7.1.3 | Clear Commands..... | 857 |
| 6.7.2 | Command Descriptions | 858 |
| 6.7.2.1 | Show Commands | 858 |
| 6.7.2.2 | Tools Commands | 866 |
| 6.7.2.3 | Clear Commands..... | 872 |
| 7 | Standards and Protocol Support | 873 |

1 Getting Started

1.1 About This Guide

This guide describes logical IP routing interfaces, virtual routers, IP and MAC-based filtering, and cflowd support and presents configuration and implementation examples.

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

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

- 7450 ESS
- 7750 SR
- 7950 XRS
- VSR

[Table 1](#) lists the available chassis types for each SR OS router.

Table 1 Supported SR OS Router Chassis Types

| 7450 ESS | 7750 SR | 7950 XRS |
|---|---|--|
| <ul style="list-style-type: none"> • 7450 ESS-7/12 running in standard mode (not mixed-mode) | <ul style="list-style-type: none"> • 7450 ESS-7/12 running in mixed-mode (not standard mode) • 7750 SR-a4/a8 • 7750 SR-c4/c12 • 7750 SR-1e/2e/3e • 7750 SR-7/12 • 7750 SR-12e | <ul style="list-style-type: none"> • 7950 XRS-16c • 7950 XRS-20/40 |

For a list of unsupported features by platform and chassis, refer to *SR OS R15.0.Rx* Software Release Notes, part number 3HE 12060 000x TQZZA or the *VSR Release Notes*, part number 3HE 12092 000x TQZZA.

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



Note: This guide generically covers Release 15.0.Rx content and may contain some content that will be released in later maintenance loads. Refer to *SR OS R15.0.Rx* Software Release Notes, part number 3HE 12060 000x TQZZA or the *VSR Release Notes*, part number 3HE 12092 000x TQZZA, for information on features supported in each load of the Release 15.0.Rx software.

1.2 Router Configuration Process

[Table 2](#) lists the tasks necessary to configure logical IP routing interfaces, virtual routers, IP and MAC-based filtering, and cflowd.

This guide is presented in an overall logical configuration flow. Each section describes a software area and provides CLI syntax and command usage to configure parameters for a functional area.

Table 2 Configuration Process

| Area | Task | Section |
|------------------------|---|---|
| Router configuration | Configure IP router parameters | Configuring IP Router Parameters |
| | Configure IP router | Configuring an IP Router with CLI |
| | Perform service management | Service Management Tasks |
| Protocol configuration | Configure VRRP parameters | Basic VRRP Configurations |
| | Configure VRRP | Common Configuration Tasks |
| | Configure VRRP policy components | Configuring VRRP Policy Components |
| | VRRP configuration management | VRRP Configuration Management Tasks |
| | Configure IP, MAC, and IP exception filter policies | Common Configuration Tasks |
| | Filter management | Filter Management Tasks |
| | Configure cFlowd | Configuring Cflowd with CLI |
| | cFlowd configuration management | Cflowd Configuration Management Tasks |

2 IP Router Configuration

2.1 Configuring IP Router Parameters

To provision services on a Nokia router, logical IP routing interfaces must be configured to associate attributes such as an IP address, port, or the system with the IP interface.

A special type of IP interface is the system interface. A system interface must have an IP address with a 32-bit subnet mask. The system interface is used as the router identifier by higher-level protocols such as OSPF and BGP, unless overwritten by an explicit router ID.

The following router features can be configured:

- [Interfaces](#)
- [Creating an IP Address Range](#)
- [Autonomous Systems \(AS\)](#)
- [Confederations](#)
- [Proxy ARP](#)

Refer to the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide* for information about DHCP and support as well as configuration examples for the 7750 SR and 7450 ESS.

2.1.1 Interfaces

Nokia routers use different types of interfaces for various functions. Interfaces must be configured with parameters such as the interface type (network and system) and address. A port is not associated with a system interface. An interface can be associated with the system (loopback address).

2.1.1.1 Network Interface

A network interface (a logical IP routing interface) can be configured on one of the following entities:

- Physical or logical port
- A SONET/SDH channel for the 7750 SR or 7450 ESS

2.1.1.2 Network Domains

To determine which network ports (and, therefore, which network complexes) are eligible to transport traffic of individual SDPs, network-domain is provided. Network-domain information is then used for the sap-ingress queue allocation algorithm applied to VPLS SAPs. This algorithm is optimized in so that no sap-ingress queues are allocated if the specified port does not belong to the network-domain used in the specified VPLS. Also, sap-ingress queues will not be allocated toward network ports (regardless of the network-domain membership) if the specified VPLS does not contain any SDPs.

Sap-ingress queue allocation considers the following:

- SHG membership of individual SDPs
- Network-domain definition under SDP to restrict the topology in which the specified SDP can be set-up

The implementation supports four network-domains within any VPLS.

Network-domain configuration at the SDP level is ignored when the SDP is used for Epipe, Lpipe, or Apipe bindings.

Network-domain configuration is irrelevant for Layer 3 services (Layer 3 VPN and/or IES service). Network-domain configuration can be defined in the base routing context and associated only with network interfaces in this context. Network domains are not applicable to loopback and system interfaces.

The network-domain information will only be used for ingress VPLS sap queue-allocation. It will not be considered by routing during SDP setup. Therefore, if the specified SDP is routed through network interfaces that are not part of the configured network domain, the packets will be still forwarded, but their QoS and queuing behavior will be based on default settings. Also, the packet will not appear in SAP statistics.

There will always be one network-domain with reserved name default. The interfaces will always belong to a default network-domain. It will be possible to assign a specific interface to different user-defined network-domains. The loopback and system interfaces will be also associated with the default network-domain at the creation. However, any attempt to associate those interfaces with any explicitly defined network-domain will be blocked at the CLI level because there is no benefit for that association.

Any SDP can be assigned only to one network domain. If none is specified, the system will assign the default network-domain. This means that all SAPs in VPLS will have queue reaching all fwd-complexes serving interfaces that belong to the same network-domains as the SDPs.

It is possible to assign/remove network-domain association of the interface/SDP without requiring deletion of the respective object.

2.1.1.3 System Interface

The system interface is associated with a network entity (such as a specific router or switch), not a specific interface. The system interface is also referred to as the loopback address. The system interface is associated during the configuration of the following entities:

- Termination point of service tunnels
- Hops when configuring MPLS paths and LSPs
- Addresses on a target router for BGP and LDP peering

The system interface is used to preserve connectivity (when routing reconvergence is possible) when an interface fails or is removed. The system interface is used as the router identifier, and a system interface must have an IP address with a 32-bit subnet mask.

2.1.1.4 Unicast Reverse Path Forwarding Check (uRPF)

uRPF helps to mitigate problems that are caused by the introduction of malformed or forged (spoofed) IP source addresses into a network by discarding IP packets that lack a verifiable IP source address. For example, a number of common types of denial-of-service (DoS) attacks, including smurf and tribe flood network (TFN), can take advantage of forged or rapidly changing source addresses to allow attackers to thwart efforts to locate or filter the attacks. For Internet service providers (ISPs) that provide public access, uRPF deflects such attacks by forwarding only packets with source addresses that are valid and consistent with the IP routing table. This action protects the network of the ISP, its customer, and the rest of the Internet.

uRPF is supported for both IPv4 and IPv6 on network and access. It is supported on any IP interface, including base router, IES, VPRN, and subscriber group interfaces.

In strict mode, uRPF checks whether the incoming packet has a source address that matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix.

In loose mode, uRPF checks whether the incoming packet has a source address that matches a prefix in the routing table; loose mode does not check whether the interface expects to receive a packet with a specific source address prefix.

Loose mode uRPF check is supported for ECMP, IGP shortcuts, and VPRN MP-BGP routes. Packets coming from a source that matches any ECMP, IGP shortcut, or VPRN MP-BGP route will pass the uRPF check even when uRPF is set to strict mode on the incoming interface.

In the case of ECMP, this allows a packet received on an IP interface configured in strict uRPF mode to be forwarded if the source address of the packet matches an ECMP route, even if the IP interface is not a next-hop of the ECMP route or not a member of any ECMP routes. The strict-no-ecmp uRPF mode may be configured on any interface that is known to not be a next-hop of any ECMP route. When a packet is received on this interface, and the source address matches an ECMP route, the packet is dropped by uRPF.

If there is a default route, the following is included in the uRPF check:

- A loose mode uRPF check always succeeds.
- A strict mode uRPF check only succeeds if the source address matches any route (including the default route) where the next-hop is on the incoming interface for the packet.

Otherwise, the uRPF check fails.

If the source IP address matches a discard/blackhole route, the packet is treated as if it failed the uRPF check.

2.1.1.5 Creating an IP Address Range

An IP address range can be reserved for exclusive use for services by defining the **config>router>service-prefix** command. When the service is configured, the IP address must be in the range specified by a service prefix. If no service prefix is configured, no limitation exists.

Addresses in the range of a service prefix can be allocated to a network port unless the exclusive parameter is specified. Then, the address range is exclusively reserved for services.

When defining a range that is a superset of a previously defined service prefix, the subset will be replaced with the superset definition. For example, if a service prefix exists for 10.10.10.0/24, and a new service prefix is configured as 10.10.0.0/16, then the old address (10.10.10.0/24) will be replaced with the new address (10.10.0.0/16).

When defining a range that is a subset of a previously defined service prefix, the subset will replace the existing superset, providing that addresses used by services are not affected. For example, if a service prefix exists for 10.10.0.0/16, and a new service prefix is configured as 10.10.10.0/24, then the 10.10.0.0/16 address will be removed, provided that no services are configured that use 10.10.x.x addresses other than 10.10.10.x.

2.1.1.6 QoS Policy Propagation Using BGP (QPPB)

This section describes QPPB as it applies to VPRN, IES, and router interfaces. Refer to the “Internet Enhanced Service” section in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* and the “IP Router Configuration” section in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Router Configuration Guide*.

QoS policy propagation using BGP (QPPB) is a feature that allows a route to be installed in the routing table with a forwarding-class and priority so that packets matching the route can receive the associated QoS. The forwarding-class and priority associated with a BGP route are set using BGP import route policies. This feature is called QPPB, even though the feature name refers to BGP specifically. On SR OS, QPPB is supported for BGP (IPv4, IPv6, VPN-IPv4, VPN-IPv6), RIP, and static routes.

SAP ingress and network QoS policies can achieve the same result as QPPB (for example, by assigning a packet arriving on an IP interface to a specific forwarding-class and priority/profile, based on the source address or destination address of the packet). However, the effort involved in creating the QoS policies, keeping them up-to-date, and applying them across many nodes is much greater than with QPPB. In a typical application of QPPB, a BGP route is advertised with a BGP community attribute that conveys a specific QoS. Routers that receive the advertisement accept the route into their routing table and set the forwarding-class and priority of the route from the community attribute.

2.1.1.6.1 QPPB Applications

There are two typical applications of QPPB:

1. Coordination of QoS policies between different administrative domains
2. Traffic differentiation within a single domain, based on route characteristics

2.1.1.6.2 Inter-AS Coordination of QoS Policies

The operator of an administrative domain “A” can use QPPB to signal to a peer administrative domain “B” that traffic sent to certain prefixes advertised by domain A should receive a specific QoS treatment in domain B. For example, an ASBR of domain A can advertise a prefix to domain B and include a BGP community attribute with the route. The community value implies a specific QoS treatment, as agreed by the two domains (in their peering agreement or service level agreement, for example). When the ASBR and other routers in domain B accept and install the route for that prefix into their routing table, they apply a QoS policy on selected interfaces that classifies traffic toward that prefix into the QoS class implied by the BGP community value.

QPPB may also be used to request that traffic sourced from specific networks receive appropriate QoS handling in downstream nodes that may span different administrative domains. This can be achieved by advertising the source prefix with a BGP community, as described. However, in this case, other approaches are equally valid, such as marking the DSCP or other CoS fields based on the source IP address, so that downstream domains can take action based on a common understanding of the QoS treatment implied by different DSCP values.

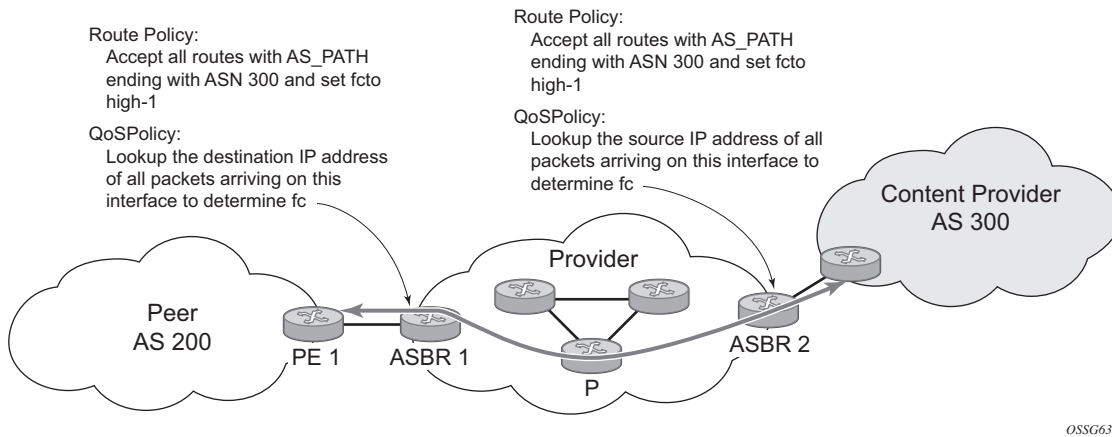
In the preceding examples, coordination of QoS policies using QPPB could be between a business customer and their IP VPN service provider, or between one service provider and another.

2.1.1.6.3 Traffic Differentiation Based on Route Characteristics

A network operator might need to provide differentiated service to specific traffic flows within its network, and these traffic flows can be identified with known routes. For example, the operator of an ISP network might need to give priority to traffic originating in a specific ASN (the ASN of a content provider offering over-the-top services to the ISP’s customers), following a specific AS_PATH, or destined for a specific next-hop (remaining on-net vs. off-net).

[Figure 1](#) shows an example of an ISP that has an agreement with the content provider managing AS300 to provide traffic sourced and terminating within AS300 with differentiated service appropriate to the content being transported. In this example, ASBR1 and ASBR2 mark the DSCP of packets terminating and sourced, respectively, in AS300 so that other nodes within the ISP’s network do not need to rely on QPPB to determine the correct forwarding-class to use for the traffic. The DSCP or other CoS markings could be left unchanged in the ISP’s network and QPPB used on every node.

Figure 1 Use of QPPB to Differentiate Traffic in an ISP Network



2.1.1.7 QPPB

There are two main aspects of the QPPB feature:

- The ability to associate a forwarding-class and priority with specific routes in the routing table.
- The ability to classify an IP packet arriving on a specific IP interface to the forwarding-class and priority associated with the route that best matches the packet.

2.1.1.7.1 Associating an FC and Priority with a Route

This feature uses the **fc** command in the route-policy hierarchy to set the forwarding class and, optionally, the priority associated with routes accepted by a route-policy entry. The command has the following structure:

```
fc fc-name [priority {low | high}]
```

The use of the **fc** command is shown by the following example:

```
config>router>policy-options
begin
community gold members 300:100
policy-statement qppb_policy
entry 10
from
protocol bgp
community gold
exit
```

```
        action accept
          fc hl priority high
        exit
      exit
    exit
  commit
```

The **fc** command is supported with all existing from and to match conditions in a route policy entry, with any action other than reject, and with next-entry, next-policy, and accept actions. If a next-entry or next-policy action results in multiple matching entries, then the last entry with a QPPB action determines the forwarding class and priority.

A route policy that includes the **fc** command in one or more entries can be used in any import or export policy, but the **fc** command has no effect except in the following types of policies:

- VRF import policies:
 - **config>service>vprn>vrf-import**
- BGP import policies:
 - **config>router>bgp>import**
 - **config>router>bgp>group>import**
 - **config>router>bgp>group>neighbor>import**
 - **config>service>vprn>bgp>import**
 - **config>service>vprn>bgp>group>import**
 - **config>service>vprn>bgp>group>neighbor>import**
- RIP import policies:
 - **config>router>rip>import**
 - **config>router>rip>group>import**
 - **config>router>rip>group>neighbor>import**
 - **config>service>vprn>rip>import**
 - **config>service>vprn>rip>group>import**
 - **config>service>vprn>rip>group>neighbor>import**

As shown, QPPB route policies support routes learned from RIP and BGP neighbors of a VPRN, as well as for routes learned from RIP and BGP neighbors of the base/global routing instance.

QPPB is supported for BGP routes belonging to any of the following address families:

- IPv4 (AFI=1, SAFI=1)
- IPv6 (AFI=2, SAFI=1)

- VPN-IPv4 (AFI=1, SAFI=128)
- VPN-IPv6 (AFI=2, SAFI=128)

A VPN-IP route may match both a VRF import policy entry and a BGP import policy entry (if `vpn-apply-import` is configured in the base router BGP instance). In this case, the VRF import policy is applied first, then the BGP import policy, so the QPPB QoS is based on the BGP import policy entry.

This feature also provides the ability to associate a forwarding-class and, optionally, priority with IPv4 and IPv6 static routes. This is achieved by specifying the forwarding-class within the **static-route-entry next-hop** or **indirect** context.

Priority is optional when specifying the forwarding class of a static route, but when configured it can only be deleted and returned to unspecified by deleting the entire static route.

2.1.1.7.2 Displaying QoS Information Associated with Routes

The following commands are enhanced to show the forwarding-class and priority associated with the displayed routes:

- `show router route-table`
- `show router fib`
- `show router bgp routes`
- `show router rip database`
- `show router static-route`

This feature uses a **qos** keyword with the **show>router>route-table** command. When this option is specified, the output includes an additional line per route entry that displays the forwarding class and priority of the route. If a route has no fc and priority information, the third line is blank. The following CLI shows an example:

show router route-table [family] [ip-prefix[/prefix-length]] [longer | exact] [protocol protocol-name] qos

An example output of this command is as follows:

```
A:Dut-A# show router route-table 10.1.5.0/24 qos
=====
Route Table (Router: Base)
=====
Dest Prefix                                Type   Proto   Age      Metric  Pref
      Next Hop[Interface Name]
      QoS
-----
10.1.5.0/24                               Remote  BGP     15h32m52s  0
```

```

PE1_to_PE2                                0
h1, high
-----
No. of Routes: 1
=====
A:Dut-A#

```

2.1.1.7.3 Enabling QPPB on an IP interface

To enable QoS classification of ingress IP packets on an interface based on the QoS information associated with the routes that best match the packets, configure the **qos-route-lookup** command in the IP interface. The **qos-route-lookup** command has parameters to indicate whether the QoS result is based on lookup of the source or destination IP address in every packet. There are separate **qos-route-lookup** commands for the IPv4 and IPv6 packets on an interface, which allows QPPB to be enabled for IPv4 only, IPv6 only, or both IPv4 and IPv6. Currently, QPPB based on a source IP address is not supported for IPv6 packets or for ingress subscriber management traffic on a group interface.

The **qos-route-lookup** command is supported on the following types of IP interfaces:

- base router network interfaces (**config>router>interface**)
- VPRN SAP and spoke SDP interfaces (**config>service>vprn>interface**)
- VPRN group-interfaces (**config>service>vprn>sub-if>grp-if**)
- IES SAP and spoke SDP interfaces (**config>service>ies>interface**)
- IES group-interfaces (**config>service>ies>sub-if>grp-if**)

When the **qos-route-lookup** command with the destination parameter is applied to an IP interface and the destination address of an incoming IP packet matches a route with QoS information, the packet is classified to the fc and priority associated with that route. The command overrides the FC and priority/profile determined from the SAP ingress or network QoS policy associated with the IP interface (see section 5.7 for more information). If the destination address of the incoming packet matches a route with no QoS information, the fc and priority of the packet remain as determined by the sap-ingress or network qos policy.

Similarly, when the **qos-route-lookup** command with the source parameter is applied to an IP interface and the source address of an incoming IP packet matches a route with QoS information, the packet is classified to the FC and priority associated with that route. The command overrides the FC and priority/profile determined from the SAP ingress or network QoS policy associated with the IP interface. If the source address of the incoming packet matches a route with no QoS information, the FC and priority of the packet remain as determined by the SAP ingress or network QoS policy.

Currently, QPPB is not supported for ingress MPLS traffic on network interfaces or on CsC PE'-CE' interfaces (**config>service>vprn>nw-if**).



Note: QPPB based on a source IP address is not supported for ingress subscriber management traffic on a group interface.

2.1.1.7.4 QPPB When Next-Hops are Resolved by QPPB Routes

In some cases (IP VPN inter-AS model C, Carrier Supporting Carrier, indirect static routes, and so on), an IPv4 or IPv6 packet may arrive on a QPPB-enabled interface and match a route A1 whose next-hop N1 is resolved by a route A2 with next-hop N2. Similarly, N2 is resolved by a route A3 with next-hop N3, and so on. The QPPB result is based only on the forwarding-class and priority of route A1. If A1 does not have a forwarding-class and priority association, the QoS classification is not based on QPPB, even if routes A2, A3, and so on, have forwarding-class and priority associations.

2.1.1.7.5 QPPB and Multiple Paths to a Destination

When ECMP is enabled, some routes may have multiple equal-cost next-hops in the forwarding table. When an IP packet matches such a route, the next-hop selection is typically based on a hash algorithm that tries to load balance traffic across all the next-hops while keeping all packets of a flow on the same path. The QPPB configuration model described in [Associating an FC and Priority with a Route](#) allows different QoS information to be associated with the different ECMP next-hops of a route. The forwarding-class and priority of a packet matching an ECMP route is based on the next-hop used to forward the packet.

When Edge PIC [1] is enabled, some BGP routes may have a backup next-hop in the forwarding table, as well as the one or more primary next-hops representing the equal-cost best paths allowed by the ECMP/multipath configuration. When an IP packet matches such a route, a reachable primary next-hop is selected (based on the hash result) but if all the primary next-hops are unreachable, the backup next-hop is used. The QPPB configuration model described in [Associating an FC and Priority with a Route](#) allows the forwarding-class and priority associated with the backup path to be different from the QoS characteristics of the equal-cost best paths. The forwarding class and priority of a packet forwarded on the backup path is based on the **fc** and priority of the backup route.

2.1.1.7.6 QPPB and Policy-Based Routing

When an IPv4 or IPv6 packet with destination address arrives on an interface with both QPPB and policy-based-routing enabled:

- There is no QPPB classification if the IP filter action redirects the packet to a directly connected interface, even if the destination address is matched by a route with a forwarding-class and priority.
- QPPB classification is based on the forwarding-class and priority of the route matching IP address Y if the IP filter action redirects the packet to the indirect next-hop IP address Y, even if the destination address is matched by a route with a forwarding-class and priority.

2.1.1.8 QPPB and GRT Lookup

Source-address based QPPB is not supported on any SAP or spoke SDP interface of a VPRN configured with the **grt-lookup** command.

2.1.1.8.1 QPPB Interaction with SAP Ingress QoS Policy

When QPPB is enabled on a SAP IP interface, the forwarding class of a packet may change from **fc1** (the original **fc** determined by the SAP ingress QoS policy) to **fc2**, the new **fc** determined by QPPB. In the ingress datapath, SAP ingress QoS policies are applied in the first P chip and route lookup/QPPB occurs in the second P chip. This has the following implications:

- Ingress remarking (based on profile state) is always based on the original **fc** (**fc1**) and sub-class (if defined).
- The profile state of a SAP ingress packet that matches a QPPB route depends on the configuration of **fc2** only. If the de-1-out-profile flag is enabled in **fc2**, and **fc2** is not mapped to a priority mode queue, the packet will be marked out of profile if its DE bit = 1. If the profile state of **fc2** is explicitly configured (in or out) and **fc2** is not mapped to a priority mode queue, the packet is assigned this profile state. In both cases, there is no consideration of whether **fc1** was mapped to a priority mode queue.
- The priority of a SAP ingress packet that matches a QPPB route depends on several factors. If the de-1-out-profile flag is enabled in **fc2** and the DE bit is set in the packet, priority will be low regardless of the QPPB priority or **fc2** mapping to profile mode queue, priority mode queue, or policer. If **fc2** is associated with a profile mode queue, the packet priority will be based on the explicitly configured profile state of **fc2** (in profile = high, out profile = low, undefined =

high), regardless of the QPPB priority or **fc1** configuration. If **fc2** is associated with a priority mode queue or policer, the packet priority will be based on QPPB (unless DE=1). If no priority information is associated with the route, the packet priority will be based on the configuration of **fc1**. If **fc1** mapped to a priority mode queue, the priority is based on DSCP/IP prec/802.1p. If **fc1** mapped to a profile mode queue, the priority is based on the profile state of **fc1**.

Table 3 summarizes these interactions.

Table 3 QPPB Interactions with SAP Ingress QoS

| Original FC object mapping | New FC object mapping | Profile | Priority (drop preference) | DE=1 override | In/out of profile marking |
|----------------------------|-----------------------|--|---|------------------|--------------------------------|
| Profile mode queue | Profile mode queue | From new base FC unless overridden by DE=1 | From QPPB, unless packet is marked in or out of profile in which case follows profile. Default: high priority. | From new base FC | From original FC and sub-class |
| Priority mode queue | Priority mode queue | Ignored | If DE=1 override then low otherwise from QPPB. If no DEI or QPPB overrides then from original dot1p/exp/DSCP mapping or policy default. | From new base FC | From original FC and sub-class |
| Policer | Policer | From new base FC unless overridden by DE=1 | If DE=1 override then low otherwise from QPPB. If no DEI or QPPB overrides then from original dot1p/exp/DSCP mapping or policy default. | From new base FC | From original FC and sub-class |
| Priority mode queue | Policer | From new base FC unless overridden by DE=1 | If DE=1 override then low otherwise from QPPB. If no DEI or QPPB overrides then from original dot1p/exp/DSCP mapping or policy default. | From new base FC | From original FC and sub-class |

Table 3 QPPB Interactions with SAP Ingress QoS (Continued)

| Original FC object mapping | New FC object mapping | Profile | Priority (drop preference) | DE=1 override | In/out of profile marking |
|----------------------------|-----------------------|--|---|------------------|--------------------------------|
| Policer | Priority mode queue | Ignored | If DE=1 override then low otherwise from QPPB. If no DEI or QPPB overrides then from original dot1p/exp/DSCP mapping or policy default. | From new base FC | From original FC and sub-class |
| Profile mode queue | Priority mode queue | Ignored | If DE=1 override then low otherwise from QPPB. If no DEI or QPPB overrides then follows original FC's profile mode rules. | From new base FC | From original FC and sub-class |
| Priority mode queue | Profile mode queue | From new base FC unless overridden by DE=1 | From QPPB, unless packet is marked in or out of profile in which case follows profile. Default: high priority. | From new base FC | From original FC and sub-class |
| Profile mode queue | Policer | From new base FC unless overridden by DE=1 | If DE=1 override then low otherwise from QPPB. If no DEI or QPPB overrides then follows original FC's profile mode rules. | From new base FC | From original FC and sub-class |
| Policer | Profile mode queue | From new base FC unless overridden by DE=1 | From QPPB, unless packet is marked in or out of profile in which case follows profile. Default: high priority. | From new base FC | From original FC and sub-class |

2.1.2 Router ID

The router ID, a 32-bit number, uniquely identifies the router within an autonomous system (AS) (see [Autonomous Systems \(AS\)](#)). In protocols such as OSPF, routing information is exchanged between areas—groups of networks that share routing information. It can be set to be the same as the loopback address. The router ID is used by both OSPF and BGP routing protocols in the routing table manager instance.

There are several ways to obtain the router ID. On each router, the router ID can be obtained in the following ways.

- Define the value in the **config>router** *router-id* context. The value becomes the router ID.
- Configure the system interface with an IP address in the **config>router>interface** *ip-int-name* context. If the router ID is not manually configured in the **config>router** *router-id* context, the system interface acts as the router ID.
- If neither the system interface or router ID are implicitly specified, the router ID is inherited from the last four bytes of the MAC address.
- The router can be obtained from the protocol level; for example, BGP.

2.1.3 Autonomous Systems (AS)

Networks can be grouped into areas. An area is a collection of network segments within an AS that have been administratively assigned to the same group. An area's topology is concealed from the rest of the AS, which results in a significant reduction in routing traffic.

Routing in the AS takes place on two levels, depending on whether the source and destination of a packet reside in the same area (intra-area routing) or different areas (inter-area routing). In intra-area routing, the packet is routed solely on information obtained within the area; no routing information obtained from outside the area can be used. This protects intra-area routing from the injection of bad routing information.

Routers that belong to more than one area are called area border routers. All routers in an AS do not have an identical topological database. An area border router has a separate topological database for each area it is connected to. Two routers, which are not area border routers, belonging to the same area, have identical area topological databases.

Autonomous systems share routing information, such as routes to each destination and information about the route or AS path, with other ASs using BGP. Routing tables contain lists of next hops, reachable addresses, and associated path cost metrics to each router. BGP uses the information and path attributes to compile a network topology.

2.1.4 Confederations

Configuring confederations is optional and should only be implemented to reduce the IBGP mesh inside an AS. An AS can be logically divided into smaller groupings called sub-confederations and then assigned a confederation ID (similar to an autonomous system number). Each sub-confederation has fully meshed IBGP and connections to other ASs outside of the confederation.

The sub-confederations have EBGP-type peers to other sub-confederations within the confederation. They exchange routing information as if they were using IBGP. Parameter values such as next hop, metric, and local preference settings are preserved. The confederation appears and behaves like a single AS.

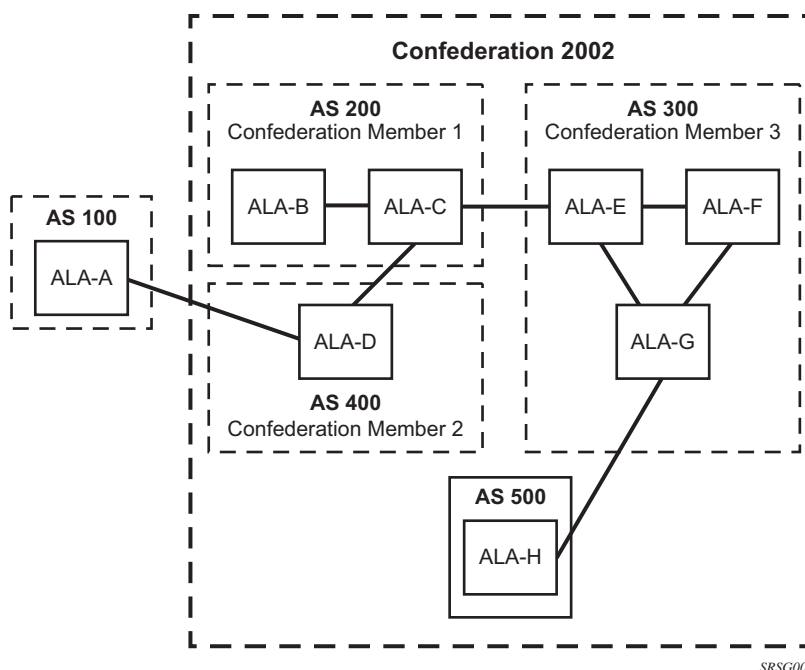
Confederations have the following characteristics:

- A large AS can be sub-divided into sub-confederations.
- Routing within each sub-confederation is accomplished via IBGP.
- EBGP is used to communicate between sub-confederations.
- BGP speakers within a sub-confederation must be fully meshed.
- Each sub-confederation (member) of the confederation has a different AS number. The AS numbers used are typically in the private AS range of 64512 to 65535.

To migrate from a non-confederation configuration to a confederation configuration requires a major topology change and configuration modifications on each participating router. Setting BGP policies to select an optimal path through a confederation requires other BGP modifications.

There are no default confederations. Router confederations must be explicitly created. [Figure 2](#) shows an example of a confederation configuration.

Figure 2 Confederation Configuration



2.1.5 Proxy ARP

Proxy ARP is the technique in which a router answers ARP requests intended for another node. The router appears to be present on the same network as the “real” node that is the target of the ARP and takes responsibility for routing packets to the “real” destination. Proxy ARP can help nodes on a subnet reach remote subnets without configuring routing or a default gateway.

Typical routers only support proxy ARP for directly attached networks; the router is targeted to support proxy ARP for all known networks in the routing instance where the virtual interface proxy ARP is configured.

To support DSLAM and other edge-like environments, proxy ARP supports policies that allow the provider to configure prefix lists that determine for which target networks proxy ARP will be attempted and prefix lists that determine for which source hosts proxy ARP will be attempted.

Also, the proxy ARP implementation will support the ability to respond for other hosts within the local subnet domain. This is needed in environments such as DSL where multiple hosts are in the same subnet but can not reach each other directly.

Static ARP is used when a Nokia router needs to know about a device on an interface that cannot or does not respond to ARP requests. The configuration can state that, if it has a packet with a specific IP address, to send it to the corresponding ARP address. Use proxy ARP so the router responds to ARP requests on behalf of another device.

2.1.6 Exporting an Inactive BGP Route from a VPRN

The **export-inactive-bgp** command under **config>service>vprn** provides an IP VPN configuration option that allows the best BGP route learned by a VPRN to be exported as a VPN-IP route even when that BGP route is inactive due to the presence of a more preferred BGP-VPN route from another PE. This “best-external” type of route advertisement is useful in active/standby multi-homing scenarios because it can ensure that all PEs have knowledge of the backup path provided by the standby PE.

2.1.7 DHCP Relay

Refer to the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide* for information about DHCP relay and support, as well as configuration examples.

2.1.8 Internet Protocol Versions

The TiMOS implements IP routing functionality, providing support for IP version 4 (IPv4) and IP version 6 (IPv6). IP version 6 (RFC 1883, *Internet Protocol, Version 6 (IPv6)*) is a version of the Internet Protocol designed as a successor to IP version 4 (IPv4) (RFC-791, *Internet Protocol*). The changes from IPv4 to IPv6 affect the following categories:

- Expanded addressing capabilities — IPv6 increases the IP address size from 32 bits (IPv4) to 128 bits, to support more levels of addressing hierarchy, a much greater number of addressable nodes, and simpler auto-configuration of addresses. The scalability of multicast routing is improved by adding a scope field to multicast addresses. Also, a type of address called an anycast address is defined that is used to send a packet to any one of a group of nodes.

- Header format simplification — Some IPv4 header fields have been dropped or made optional to reduce the common-case processing cost of packet handling and to limit the bandwidth cost of the IPv6 header.
- Improved support for extensions and options — Changes in the way IP header options are encoded allows for more efficient forwarding, less stringent limits on the length of options, and greater flexibility for introducing options in the future.
- Flow labeling capability — The capability to enable the labeling of packets belonging to traffic flows for which the sender requests special handling, such as non-default quality of service or “real-time” service was added in IPv6.
- Authentication and privacy capabilities — Extensions to support authentication, data integrity, and (optional) data confidentiality are specified for IPv6.

Figure 3 IPv6 Header Format



al_0892

Table 4 IPv6 Header Field Descriptions

| Field | Description |
|----------------|---|
| Version | 4-bit Internet Protocol version number = 6. |
| Prio. | 4-bit priority value. |
| Flow Label | 24-bit flow label. |
| Payload Length | 16-bit unsigned integer. The length of payload, for example, the rest of the packet following the IPv6 header, in octets. If the value is zero, the payload length is carried in a jumbo payload hop-by-hop option. |

Table 4 IPv6 Header Field Descriptions (Continued)

| Field | Description |
|---------------------|--|
| Next Header | 8-bit selector. Identifies the type of header immediately following the IPv6 header. This field uses the same values as the IPv4 protocol field. |
| Hop Limit | 8-bit unsigned integer. Decrement by 1 by each node that forwards the packet. The packet is discarded if the hop limit is decremented to zero. |
| Source Address | 128-bit address of the originator of the packet. |
| Destination Address | 128-bit address of the intended recipient of the packet (possibly not the ultimate recipient if a routing header is present). |

2.1.8.1 IPv6 Address Format

IPv6 uses a 128-bit address, as opposed to the IPv4 32-bit address. Unlike IPv4 addresses, which use the dotted-decimal format, with each octet assigned a decimal value from 0 to 255, IPv6 addresses use the colon-hexadecimal format X:X:X:X:X:X:X:X, where each X is a 16-bit section of the 128-bit address. For example:

2001:0DB8:0000:0000:0000:0000:0000:0000

Leading zeros must be omitted from each block in the address. A series of zeros can be replaced with a double colon. For example:

2001:DB8::

The double colon can only be used once in an address.

The IPv6 prefix is the part of the IPv6 address that represents the network identifier, which appears at the beginning of the address. The IPv6 prefix length, which begins with a forward slash (/), shows how many bits of the address make up the network identifier. For example, the address 1080:6809:8086:6502::1/64 means that the first 64 bits of the address represent the network identifier; the remaining 64 bits represent the node identifier.




Note: In SR OS 12.0.R4 and later, any function that displays an IPv6 address or prefix changes to reflect rules described in RFC 5952, *A Recommendation for IPv6 Address Text Representation*. Specifically, hexadecimal letters in IPv6 addresses are now represented in lowercase, and the correct compression of all leading zeros is displayed. This changes visible display output compared to previous SR OS releases. Previous SR OS behavior can cause issues with operator scripts that use standard IPv6 address expressions and with libraries that have standard IPv6 parsing as per RFC 5952 rules.

2.1.8.2 IPv6 Applications

Examples of the IPv6 applications supported by the TiMOS include:


- IPv6 Internet exchange peering — [Figure 4](#) shows an IPv6 Internet exchange where multiple ISPs peer over native IPv6.

Figure 4 IPv6 Internet Exchange




- IPv6 transit services — [Figure 5](#) shows IPv6 transit services provided by an ISP.

Figure 5 IPv6 Transit Services



- IPv6 services to enterprise customers and home users — [Figure 6](#) shows IPv6 services to enterprise and home broadband users.

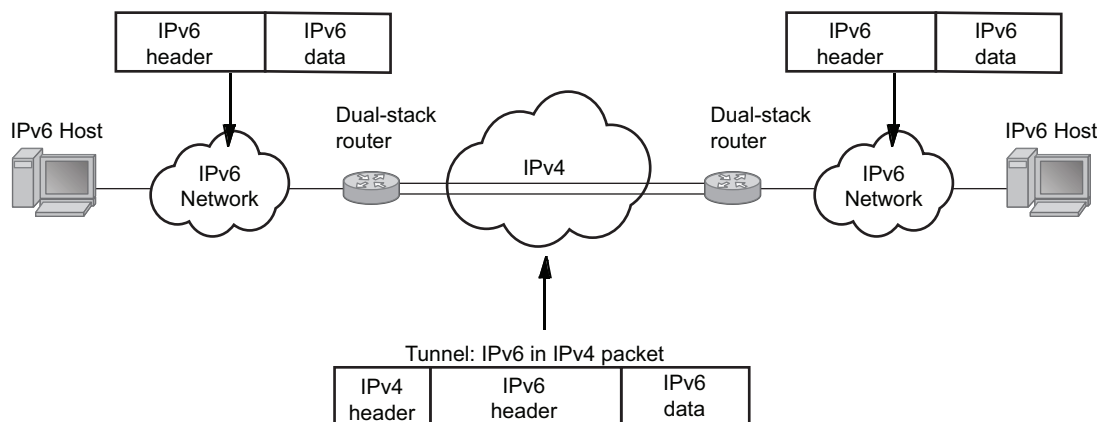
Figure 6 IPv6 Services to Enterprise Customers and Home Users

IPIPE_009

- IPv6 over IPv4 relay services — IPv6 over IPv4 tunnels are one of many IPv6 transition methods to support IPv6 in an environment where not only IPv4 exists but native IPv6 networks depend on IPv4 for greater IPv6 connectivity. Nokia routers support dynamic IPv6 over IPv4 tunneling. The IPv4 source and destination address are taken from configuration, the source address is the IPv4 system address and the IPv4 destination is the next hop from the configured IPv6 over IPv4 tunnel.

IPv6 over IPv4 is an automatic tunnel method that gives a prefix to the attached IPv6 network. [Figure 7](#) shows IPv6 over IPv4 tunneling to transition from IPv4 to IPv6.

Figure 7 IPv6 over IPv4 Tunnels



Fig_29a

2.1.8.3 DNS

The DNS client is extended to use IPv6 as transport and to handle the IPv6 address in the DNS AAAA resource record from an IPv4 or IPv6 DNS server. An assigned name can be used instead of an IPv6 address because IPv6 addresses are more difficult to remember than IPv4 addresses.

2.1.8.4 Secure Neighbor Discovery (SeND)

Secure Neighbor Discovery (SeND) in conjunction with Cryptographically Generated Addresses (CGAs) allows operators to secure IPv6 neighbor discovery between nodes on a common Layer 2 network segment.

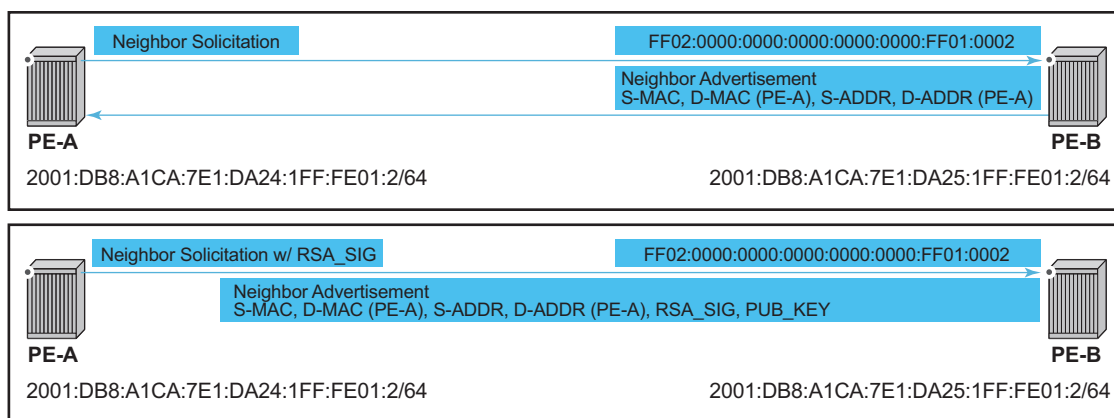
When SeND is enabled on an interface, CGAs must be enabled and static GUA/LLA IPv6 addressing is not supported. In this case, the router will generate a CGA from the configured prefix (GUA, LLA) and use that address for all communication. The router will validate NS/ND messages from other nodes on the network segment, and only install them in the neighbor cache if they pass validation.

A number of potential use-cases for SeND exist in order to secure the network from deliberate or accidental tampering during neighbor discovery, SeND can prevent hijacking of in-use IPv6 addressing or man-in-the-middle attacks, but also to validate whether a node is permitted to participate in neighbor discovery, or validate which routers are permitted to act as default gateways.

SeND affects the following areas of neighbor discovery:

- Neighbor solicitation (solicited-node multicast address; target address)
- Neighbor advertisement (solicited; unsolicited)
- Router solicitation
- Router advertisement
- Redirect messages

Figure 8 Neighbor discovery with and without SeND



al_0747

When SeND is enabled on a node, basic neighbor discovery messaging is changed as shown in [Figure 8](#). In the example, PE-A needs to find the MAC address of PE-B.

1. PE-A sends an NS message to the solicited node multicast address for PE-B's address with the CGA option, RSA signature option, timestamp option, and nonce option.
2. PE-B processes the NS message and, because it is configured for SeND operation, processes the NS. PE-B will validate the source address of the packet to ensure it is a valid CGA, then validate the cryptographic signature embedded in the NS message.
3. PE-B generates an NA message, which is sent back to PE-A with the solicited bit, router bit set. The source address is that of PE-B, while the destination address is that of PE-A from the NS message. The timestamp is generated from PE-B, while the nonce is copied from PE-A's NS message.

4. PE-A receives the NA and completes similar checks as PE-B did.

If all steps process correctly, both nodes will install each other's addresses into their neighbor cache database.

2.1.8.5 SeND Persistent CGAs

Persistent CGAs is a feature of SeND.

Previously, all generated CGAs on SeND-enabled interfaces remained unchanged after a CPM switchover, but after a reboot from a saved configuration file, all CGAs were regenerated.

To keep the same CGAs after a reboot from a saved configuration file:

1. Save the RSA key pair used for SeND.
2. Save the modifiers used during the CGA generation.

To make the CGAs persistent:

1. Import an online or offline generated RSA key pair for SeND.
2. Ensure that the CompactFlash (CF) files containing an RSA key pair that is used for SeND, are synchronized to the standby CPM by making use of the HA infrastructure used for certificates.
3. Ensure that the configuration file is saved when one or more CGAs are generated.

2.1.8.5.1 Persistent RSA Key Pair

The RSA key pair is stored in a file on the CF.

Generate an RSA Key Pair

To generate an RSA key pair, use the **admin certificate gen-keypair** command:

admin certificate gen-keypair *local-url* [**type rsa**] **size 1024**

For example:

```
admin certificate gen-keypair cf1:\myDir\myRsaKeyPair type rsa size 1024
```

This generates a der formatted file.

Import an online/offline generated RSA key pair

To import a generated RSA key pair, use the **admin certificate secure-nd-import** command:

admin certificate secure-nd-import *local-url* **format** {der | pem | pkcs12}
[password <password>] [key-rollover]

For example:

```
admin certificate secure-nd-import cfl:\myDir\myRsaKeyPair format der
```

- Because SeND only uses RSA key pairs, the command is refused if the imported key type is not RSA.
- Because SeND only supports key size 1024, the command is refused if the imported key size is not 1024.
- The password has to be specified when an offline generated file in pkcs12 format has to be imported.
- **key-rollover** keyword: see the *RSA key pair rollover mechanism* section that follows.
- This command creates the file cfx:\system-pki\secureNdKey (fixed directory and file name) and saves the imported key in that file in encrypted der format (same as the **admin certificate import** command).
- The RSA key pair is uploaded in the memory of SeND.

RSA key pair rollover mechanism

To trigger a key rollover, use the **admin certificate secure-nd-import** command described in the previous section [Import an online/offline generated RSA key pair](#).

For example:

```
admin certificate secure-nd-import cfl:\myDir\myOtherRsaKeyPair format der key-rollover
```

- If CGAs exist that are generated based on an auto-generated or previously imported RSA key pair and the **key-rollover** keyword is not specified, the **secure-nd-import** command is refused.
- If a **secure-nd-import** with **key-rollover** is requested while a previous key rollover is still being handled, the new command is refused.
- If the **secure-nd-import** command is accepted, the imported RSA key pair is written to the file cfx:\system-pki\secureNdKey and loaded to SeND. Existing CGAs if any will be regenerated.

- While handling a key rollover, SeND keeps track of which interface uses which RSA key pair. Temporarily, SeND can have two RSA key pairs in use. At all times, only the latest RSA key pair is stored in the file `cfx:\system-pki\secureNdKey`. When the rollover is finished, the RSA key pair that is no longer referred to, is deleted from SeND's memory.

Auto-generation of RSA key pair

The first time an interface becomes SeND enabled, SeND needs an RSA key pair to generate or check a modifier and to generate a CGA.

If the operator did not import an RSA key pair for SeND, an auto-generated RSA key pair will be used as a fallback.

The auto-generated RSA key pair is synchronized to the standby CPM, but will not be written to the CF. Therefore, all CGAs generated via an auto-generated RSA key pair are not persistent. A warning will be raised whenever a non-persistent CGA is generated.

The **admin certificate secure-nd-import** command without the **key-rollover** keyword will be refused if CGAs exist that made use of the auto-generated RSA key pair. Specifying the **key-rollover** keyword will result in regeneration of the CGAs.

See the section [Making non-persistent CGAs persistent](#) for more information about the procedure to make non-persistent CGAs persistent.

HA

For the synchronization of the RSA key pair file in `cfx:\system-pki\` used by SeND, the following commands for manual and automatic certificate synchronization are used:

- manual: **admin redundancy synchronize cert**
- automatic: **configure redundancy cert-sync**

SeND also synchronizes the RSA key pair to the standby CPM.

2.1.8.5.2 Persistent CGA Modifier

The modifier used during the CGA generation will be saved in the configuration file. The CGA itself is not stored.

Based on the stored modifier and RSA key pair, the same CGA can be regenerated.

The modifier is needed to be sent out in ND messages.

By storing the modifier in the configuration file, the operator can also configure an offline generated modifier (possibly with a security parameter > 1).

Example 1: Configure a SeND interface without modifiers:

```
configure router interface itf1
  address 10.10.10.1
  port 1/1/1
  ipv6
    secure-nd
    no shutdown
```

=> A modifier is generated based on the actual RSA key pair (that is, imported or auto-generated). The modifier is used to generate a link-local CGA.

=> The modifier is saved in the interface configuration file:

```
exit
address 2000:1::/64
```

=> A modifier is generated based on the actual RSA key pair. The modifier is used to generate the global CGA.

=> The modifier is stored in the interface configuration file.

Example 2: Configure a SeND interface with modifiers:

```
configure router interface itf2
  address 10.10.10.2
  port 1/1/2
  ipv6
    secure-nd
    link-local-modifier 0xABCD
```

=> The offline generated modifier is used to generate the link-local CGA:

```
no shutdown
exit
address 3000:1::/64
```

=> A modifier is generated based on the actual RSA key pair. The modifier is used to generate the global CGA.

=> The modifier is stored in the interface configuration file:

```
address 3000:2::/64 modifier 0xABCD
```

=> The same offline generated modifier as the preceding link-local address is used for the generation of a global address:

```
address 3000:3::/64 modifier 0xABCD
```

=> Another offline generated modifier (*) is used for the generation of a global address.

=> For an offline generated modifier, a check is performed to see if it is generated with the actual RSA key pair and the security parameter applicable for the interface. If this check fails, the command is refused, unless the command is triggered in the context of an exec of a config file. In that case, the modifier is replaced by a new one that is generated based on the actual RSA key pair.

2.1.8.5.3 Making non-persistent CGAs persistent

CGAs can be non-persistent because:

- The operator forgot to configure an RSA key pair for SeND, so hence the CGAs were generated based on an auto-generated RSA key pair.
- The operator forgot to synchronize an RSA key pair file to the stand-by CPM and a switch-over happens.
- The CGAs were generated by a software version not having persistent CGAs (such as, ISSU).
- The system was booted from a configuration file generated by a software version not having persistent CGAs.

Key rollover

You can import a new RSA key pair for SeND with the **key-rollover** keyword. This will result in the regeneration of all CGAs on all interfaces.

Exporting the SeND RSA key pair

Another method that does not result in the regeneration of the CGAs is to export the RSA key pair that is currently in use by SeND to the system-pki directory via an admin command:

admin certificate secure-nd-export

This command will write the RSA key pair to the file cfx:\system-pki\secureNdKey in encrypted der format.

2.1.8.5.4 Booting from a saved configuration file

Configuration saved by a software version with persistent CGAs

The file cfx:\system-pki\secureNdKey should exist. This file will be automatically uploaded by SeND during initialization.

The configuration file should contain a modifier for each address on a SeND enabled interface.

Modifiers in the configuration file are checked against the current RSA key pair. If the check fails, a new modifier and CGA is generated and a warning is raised that a new CGA is generated.

If a modifier is missing from the configuration file for an IPv6 /64 prefix on a SeND enabled interface, a new modifier and CGA will be generated based on the active RSA key pair.

Configuration saved by a software version having non-persistent CGAs


The file cfx:\system-pki\secureNdKey does not exist nor does the configuration file contain a modifier for any of the IPv6 /64 prefixes on secure-nd enabled interfaces.

New CGAs have to be generated (from the CLI context). Follow one of the procedures described in section [Making non-persistent CGAs persistent](#) to make the non-persistent CGA's persistent.

2.1.8.6 IPv6 Provider Edge Router over MPLS (6PE)

6PE allows IPv6 domains to communicate with each other over an IPv4 MPLS core network. Because forwarding is based on MPLS labels, backbone infrastructure upgrades and core router re-configuration is not required in this architecture. 6PE is a cost-effective solution for IPv6 deployment.

Figure 9 Example of a 6PE Topology within One AS



Fig_30

2.1.8.6.1 6PE Control Plane Support

The 6PE MP-BGP routers support:

- IPv4 and IPv6 dual-stack
- MP-BGP to exchange IPv6 reachability information:
 - The 6PE routers exchange IPv6 reachability information using MP-BGP (AFI 2, SAFI 4).
 - An IPv4 address of the 6PE router is encoded as an IPv4-mapped IPv6 address in the BGP next-hop field. This is usually the IPv4 system address.
 - The 6PE router binds MPLS labels to the IPv6 prefixes it advertises. SR OS routers advertise the IPv6 explicit null (value 2) in advertised 6PE routes but accept any arbitrary label from peers.
- The most preferred tunnel to the BGP next-hop allowed by the 6PE resolution filter (**config>router>bgp>next-hop-resolution>labeled-routes>transport-tunnel>family label-ipv6>resolution-filter**) is used to tunnel the traffic to the remote 6PE router.

2.1.8.6.2 6PE Data Plane Support

The ingress 6PE router can push two or more MPLS labels to send the packets to the egress 6PE router. The top labels are associated with resolving the transport tunnels. The bottom label is advertised in MP-BGP by the remote 6PE router. Typically, the IPv6 explicit null (value 2) label is used, but any arbitrary value can be received when the remote 6PE router is not an SR OS router.

The egress 6PE router pops the top transport labels. When the IPv6 explicit null label is exposed, the egress 6PE router knows that an IPv6 packet is encapsulated. It pops the IPv6 explicit null label and performs an IPv6 route lookup to find the next hop for the IPv6 packet.

2.1.9 Static Route Resolution Using Tunnels

The user can forward packets of a static route to an indirect next-hop over a tunnel programmed in TTM by configuring the following static route tunnel binding command:

```
config>router>static-route-entry {ip-prefix/prefix-length} [mcast] indirect {ip-
address}
    tunnel-next-hop
        resolution {any|disabled|filter}
        resolution-filter
            [no] ldp
            [no] rsvp-te
            [no] lsp <name1>
            [no] lsp <name2>
            .
            .
            [no] lsp <namen>
        exit
        [no] sr-isis
        [no] sr-ospf
        [no] sr-te
            [no] lsp <name1>
            [no] lsp <name2>
            .
            .
            [no] lsp <namen>
        exit
    [no] disallow-igp
    exit
exit
```

If **tunnel-next-hop** context is configured and **resolution** is set to **disabled**, the binding to tunnel is removed and resolution resumes in RTM to IP next-hops.

If **resolution** is set to **any**, any supported tunnel type in static route context will be selected following TTM preference.

The following tunnel types are supported in a static route context: LDP, RSVP-TE, Segment Routing (SR) Shortest Path, and Segment Routing Traffic Engineering (SR-TE):

- LDP

The **ldp** value instructs the code to search for an LDP LSP with a FEC prefix corresponding to the address of the indirect next-hop. Both LDP IPv4 FEC and LDP IPv6 FEC can be used as the tunnel next-hop. However, only an indirect next-hop of the same family (IPv4 or IPv6) as the prefix of the route can use an LDP FEC as the tunnel next-hop. In other words, an IPv4 (IPv6) prefix can only be resolved to an LDP IPv4 (IPv6) FEC.

- RSVP-TE

The **rsvp-te** value instructs the code to search for the set of lowest metric RSVP-TE LSPs to the address of the indirect next-hop. The LSP metric is provided by MPLS in the tunnel table. The static route treats a set of RSVP-TE LSPs with the same lowest metric as an ECMP set.

The user has the option of configuring a list of RSVP-TE LSP names to be used exclusively instead of searching in the tunnel table. In that case, all LSPs must have the same LSP metric in order for the static route to use them as an ECMP set. Otherwise, only the LSPs with the lowest common metric value are selected.

A P2P auto-lsp that is instantiated via an LSP template can be selected in TTM when **resolution** is set to **any**. However, it is not recommended to configure an auto-lsp name explicitly under the **rsvp-te** node as the auto-generated name can change if the node reboots, which will blackhole the traffic of the static route.

- SR Shortest Path

When the **sr-isis** or **sr-ospf** value is enabled, an SR tunnel to the indirect next-hop is selected in the TTM from the lowest preference ISIS or OSPF instance, and if many instances have the same lowest preference, it is selected from the lowest numbered IS-IS or OSPF instance. Both SR-ISIS IPv4 and SR-ISIS IPv6 tunnels can be used as tunnel next-hops. However, only an indirect next-hop of the same family (IPv4 or IPv6) as the prefix of the route can use an SR-ISIS tunnel as a tunnel next-hop. In other words, an IPv4 (IPv6) prefix can only be resolved to a SR-ISIS IPv4 (IPv6).

- SR-TE

The **sr-te** value instructs the code to search for the set of lowest metric SR-TE LSPs to the address of the indirect next-hop. The LSP metric is provided by MPLS in the tunnel table. The static route treats a set of SR-TE LSPs with the same lowest metric as an ECMP set.

The user has the option of configuring a list of SR-TE LSP names to be used exclusively instead of searching in the tunnel table. In that case, all LSPs must have the same LSP metric in order for the static route to use them as an ECMP set. Otherwise, only the LSPs with the lowest common metric value are selected.

If one or more explicit tunnel types are specified using the **resolution-filter** option, only these tunnel types will be selected again following the TTM preference.

The user must set **resolution** to **filter** to activate the list of tunnel-types configured under resolution-filter.

If **disallow-igp** is enabled, the static route will not be activated using IP next-hops in RTM if no tunnel next-hops are found in TTM.

2.1.9.1 Static Route ECMP Support

The following is the ECMP behavior of a static route:

- ECMP is supported when resolving in RTM multiple static routes of the same prefix with multiple user-entered indirect IP next-hops. The system picks as many direct next-hops as available in RTM beginning from the first indirect next-hop and up to the value of the **ecmp** option in the system.
- ECMP is also supported when resolving in TTM a static route to a single indirect next-hop using a LDP tunnel when LDP has multiple direct next-hops.
- ECMP is supported when resolving in TTM a static route to a single indirect next-hop using a RSVP-TE tunnel type when there is more than one RSVP LSP with the same lowest metric to the indirect next-hop.
- ECMP is supported when resolving in TTM a static route to a single indirect next-hop using a list of user-configured RSVP-TE LSP names when these LSPs have the same metric to the indirect next-hop.
- ECMP is supported when resolving in TTM multiple static routes of the same prefix with multiple user-entered indirect next-hops, each binding to a tunnel type. The system picks as many tunnel next-hops as available in TTM beginning from the first indirect next-hop and up to the value of the **ecmp** option in the system. The spraying of flow packets is performed over the entire set of resolved next-hops that correspond to the selected indirect next-hops.
- ECMP is supported when resolving concurrently in RTM and TTM multiple static routes of the same prefix with multiple user-entered indirect tunnel next-hops. There is no support for mixing IP and tunnel next-hops for the same prefix using different indirect next-hops. Tunnel next-hops are preferred over IP next-hops.

2.2 Weighted Load Balancing over MPLS LSP

The weighted load-balanced, or weighted-ecmp, feature sprays packets of IGP, BGP, and static route prefixes, resolved to a set of ECMP tunnel next hops, proportionally to the weights configured for each MPLS LSP in the ECMP set.

Weighted load balancing is supported in the following forwarding contexts:

- IGP prefix resolved to IGP shortcuts in RTM (**igp-shortcut** or **advertise-tunnel-link** enabled in the IGP instance)
- BGP prefix with the BGP next hop resolved to IGP shortcuts in RTM (**rsvp-shortcut** enabled in the IGP instance)
- static route prefix resolved to an indirect next hop, which is resolved to a set of equal-metric MPLS LSPs in TTM. The user can allow automatic selection or specify the names of the equal-metric MPLS LSPs in TTM to be used in the ECMP set.
- static route prefix resolved to an indirect next hop, which is resolved to IGP shortcuts in RTM
- BGP prefix with a BGP next hop resolved to a static route, which resolves to a set of tunnel next hops toward an indirect next hop in RTM or TTM
- BGP prefix resolving to another BGP prefix, whose next hop is resolved to a set of ECMP tunnel next hops with a static route in RTM or TTM or to IGP shortcuts in RTM
- BGP labeled IPv6 packets (6PE) over RSVP LSPs resolving in TTM

This feature does not modify the route calculation: the same set of ECMP next hops is computed for a prefix. The feature also does not change the hash routine; only the spraying of the flows over the tunnel next hops is modified to reflect the normalized weight of each tunnel next hop.

Static route implementation supports ECMP over a set of equal-cost MPLS LSPs. The user can allow automatic selection or specify the names of the equal-metric MPLS LSPs in TTM to be used in the ECMP set. For more information, see [Static Route Resolution Using Tunnels](#).

2.2.1 Weighted Load Balancing IGP, BGP, and Static Route Prefix Packets over IGP Shortcut

2.2.1.1 Feature Configuration

The user must have the IGP shortcut or forwarding adjacency feature enabled in one or more IGP instances:

```
config>router>ospf(isis)>igp-shortcut
```

```
config>router>ospf(isis)>advertise-tunnel-link
```

The user can also disable specific MPLS LSPs from being used in IGP shortcut or forwarding adjacency by configuring the following:

```
config>router>mpls>lsp>no igp-shortcut
```

The user enables the weighted load balancing feature using the following router level command:

```
config>router>weighted-ecmp
```

When this command is enabled, packets of IGP, BGP, and static route prefixes resolved to a set of ECMP tunnel next-hops are sprayed proportionally to the weights configured for each MPLS LSP in the ECMP set.

The user can configure a weight for each LSP using the following command:

```
config>router>mpls>lsp>load-balancing-weight <32-bit-integer>
```

For an auto-LSP signaled via an LSP template, the weight is configured using the following command:

```
config>router>mpls>lsp-template>load-balancing-weight <32-bit-integer>
```

There is no default weight value for an LSP. If any LSP in the ECMP set of a prefix does not have a weight configured, the regular ECMP spraying for the prefix will be performed. The user-entered weight is normalized to the closest integer value that represents the number of entries in the ingress prefix hash table assigned to the LSP for the purpose of spraying packets of all prefixes resolved to this LSP. The higher the normalized weight, the more entries will be assigned to the LSP, the more packets will be sent to this LSP.

2.2.1.2 Feature Behavior

This section describes the behavior of the weighted load-balancing feature for IGP, BGP, and static route prefixes resolved in RTM to IGP shortcuts.

When an IGP, BGP, or a static route prefix is resolved in RTM to a set of ECMP tunnel next-hops of type RSVP-TE, and the router level **weighted-ecmp** option is enabled, the ingress hash table for the next-hop selection is populated with a number of tunnel next-hop entries for each LSP equal to the normalized LSP weight value. All prefixes resolving to the same set of ECMP tunnel next-hops use the same table.

This feature performs the following:

1. MPLS populates the user-configured LSP weight in TTM. When the global command **weighted-ecmp** is enabled, and any LSP in the ECMP set of a prefix does not have a weight configured, the regular ECMP spraying for the prefix will be performed.
2. IGP computes the normalized weight for each prefix tunnel next-hop. The minimum value of the normalized weight is 1 and the maximum is 64. IGP updates the route in RTM with the set of tunnel next-hops and normalized weights. RTM downloads the information to IOM for inclusion in the FIB.
3. The normalized weights of route tunnel next-hops are updated in the following cases:
 - When the main SPF is run following a trigger, for example, network failure, and updates a route with a modified set of tunnel next-hops. This will trigger a route re-download to the IOM and all users of RTM are notified.
 - The user adds or changes the weight of one or more LSPs. In this case, RTM will perform a route download to IOM, but other users of RTM are not notified because the route resolution did not change.
4. The weighted load balancing feature is only applied to a prefix when all the tunnel next-hops in the ECMP set have the same endpoint. If an IGP prefix resolves in RTM to a set of ECMP tunnel next-hops that do not terminate on the same endpoint, the regular ECMP spraying is performed. If BGP performs BGP ECMP to a set of BGP ECMP next-hops for a prefix (**weighted-bgp-ecmp-prd**), regular ECMP spraying is performed toward a BGP next-hop if the subset of its tunnel next-hops does not terminate on the same endpoint.
5. Regular ECMP spraying is also applied if a prefix is resolved in RTM to an ECMP set that consists of a mix of IP and tunnel next-hops.
6. This feature is not supported in the following contexts:
 - Packets of BGP prefix with the BGP next-hop resolved in TTM to RSVP LSP (BGP shortcut).

- CPM generated packets, including OAM packets, which are looked-up in RTM and which are forwarded over tunnel next-hops. These will be forwarded using either regular ECMP or by selecting one next-hop from the set.

2.2.1.3 ECMP Considerations

The weight assigned to an LSP affects only the forwarding decision, not the routing decision. It does not change the selection of the set of ECMP tunnel next-hops of a prefix when more next-hops exist than the value of the router **ecmp** option. This selection continues to follow the algorithm used in the IGP shortcut feature.

After the set of tunnel next-hops is selected, the LSP weight is used to modulate the amount of packets forwarded over each next-hop.

2.2.1.4 Weighted Load Balancing Static Route Packets over MPLS LSP

2.2.1.4.1 Feature Configuration

The configuration of the resolution of a static route prefix to set of MPLS LSPs is described in [Static Route Resolution Using Tunnels](#) which also provides the selection rules among multiple LSP types: RSVP-TE, SR-TE, LDP, SR-ISIS, and SR-OSPF. A static route of a prefix can only be resolved to a set of tunnel next-hops of the same type though, for each indirect next-hop.

To perform ECMP over a set of configured MPLS LSPs, the user must enter two or more LSP names to be used as tunnel next-hops. If automatic selection is performed, ECMP is performed if two or more MPLS LSPs are in TTM to the indirect next-hop of the static route. However, all LSPs must have the same LSP metric; otherwise, only the tunnel next-hops with the same lowest metric will be activated for the static route.

The user can force the metric of an LSP to a constant value using the following command:

CLI Syntax: `config>router>mpls>lsp>metric`

If the user enters, for the same static route, more LSP names with the same LSP metric than the value of the router level **ecmp** option, only the first configured LSPs equal to the **ecmp** value will be selected. The remaining tunnel next-hops for the route will not be activated. When automatic MPLS LSP selection is performed in TTM, the lowest tunnel ID is used as a tie-breaker among the same lowest metric LSPs.

To perform weighted load-balancing over the set of MPLS LSPs, either when the LSP names are provided or when auto-selection in TTM is performed, the user must also enable the weighted ECMP globally like for static, IGP, and BGP prefixes resolving to IGP shortcuts:

CLI Syntax: `config>router>weighted-ecmp`

2.2.1.4.2 Feature Behavior

The behavior of this feature in terms of RTM and IOM is exactly the same as in the case of BGP, IGP, and static route prefixes resolving to IGP shortcuts. See [Feature Behavior](#) for more information. In this case, the static route module computes the normalized weight for each prefix tunnel next-hop of the static route indirect next-hop. The minimum value of the normalized weight is 1 and the maximum is 64. The static route module updates the route in RTM with the set of tunnel next-hops and normalized weights. RTM downloads the information to IOM for inclusion in the FIB.

If any LSP in the ECMP set of a prefix static route does not have a weight configured, the regular ECMP spraying for the prefix will be performed.

ECMP is also supported when resolving in TTM the same static route with multiple user-entered indirect next-hops, each binding to the same or different tunnel types. The system picks as many tunnel next-hops as available in RTM, beginning from the first indirect next-hop and up to the value of the **ecmp** option in the system. In this case, the weighted load-balancing will be applied directly using the weights of the selected set of tunnel next-hops. If any LSP in the ECMP set of a prefix static route does not have a weight configured, or if any of the indirect next-hops binds to an LDP LSP, the regular ECMP spraying for the prefix will be performed.

If the same prefix is resolved via both a static route and an IGP shortcut route, the RTM default protocol preference will install the static route only. Therefore, the set of ECMP tunnel next-hops and the weighted load balancing behavior will be determined by the static route configuration and not by the IGP shortcut configuration.

2.2.2 Weighted Load Balancing for 6PE

ECMP-like spraying for BGP labeled IPv6 packets (6PE) is controlled using the **config>router>ecmp** *max-ecmp-routes* command, where *max-ecmp-routes* represents the maximum number of RSVP tunnels in the set representing equal-cost paths to the BGP next hop.

Weighted ECMP behavior, where the load-balancing weight of the RSVP tunnel is considered in the packet spraying behavior, is configured using the **config>router>bgp>next-hop-resolution>weighted-ecmp** command. Weighted ECMP is disabled by default.

2.3 Class-Based Forwarding of IPv4/IPv6 Prefix Over IGP IPv4 Shortcut

This feature enables class-based forwarding (CBF) over IGP shortcuts. When the **class-forwarding** command is enabled, the following types of packets are forwarded based on their forwarding class:

- packets of BGP prefixes
- packets that are CPM-originated for the IPv4, IPv6, or both IPv4 and IPv6 families that have been enabled over IGP shortcuts using the **igp-shortcut** CLI context in one or more IGP instances

The SR OS CBF implementation supports spraying of packets over a maximum of four forwarding sets of ECMP LSPs. The user must define a class-forwarding policy object in MPLS to configure the mapping of FCs to the forwarding sets. Then, the user assigns the CBF policy name and set ID to each MPLS LSP that is used in IGP shortcuts.

When a BGP IPv4 or IPv6 prefix is resolved, the FC of the packet, is used to look up the forwarding set ID. Then, a modulo operation is performed on the tunnel next-hops of this set ID only, to spray packets of this FC. The data path concurrently implements, CBF and ECMP within the tunnels of each set ID.

CPM-originated packets on the router, including control plane and OAM packets, are forwarded over a single LSP from the set of LSPs that the packet's FC is mapped to, as per the CBF configuration.

2.3.1 Feature Configuration

The user enables CBF over IGP shortcuts using the following command:

```
A:Reno 194# configure>router$ class-forwarding
```

```
config
  router
    [no] mpls
      class-forwarding-policy policy-name create
      fc be forwarding-set set-id <1..4>
      fc l2 forwarding-set set-id <1..4>
      fc af forwarding-set set-id <1..4>
      fc l1 forwarding-set set-id <1..4>
      fc h2 forwarding-set set-id <1..4>
      fc ef forwarding-set set-id <1..4>
      fc h1 forwarding-set set-id <1..4>
      fc nc forwarding-set set-id <1..4>
```

```
[no] default-set set-id <1..4>
```

All FCs are mapped to set 1 as soon as the policy is created. The user can make changes to the mapping of FCs as required. An FC, which is not added to the class-forwarding policy, is thus always mapped to set 1. At most, an FC can be mapped to a single forwarding set. One or more FCs can map to the same set. The user can indicate the initial default set by including the **default-set** option.

The default forwarding set is used to forward packets of any FC in cases where all LSPs of the forwarding set the FC maps to become operationally down. The router uses the user-configured default set as the initial default set. Otherwise, the router elects the lowest numbered set as the default forwarding set in a class-forwarding policy. When the last LSP in a default forwarding set goes into an operationally down state, the router designates the next lowest-numbered set as the new default forwarding set.

A mapping to a class-forwarding policy and set is added to the existing CBF configuration of an RSVP-TE LSP and to the LSP template of an RSVP-TE auto-LSP. The following commands perform this function.

```
A:Reno 194# config>router>mpls>lsp>class-forwarding$ forwarding-set policy  
policy-name set set-id <1..4>
```

```
A:Reno 194# config>router>mpls>lsp-template>class-forwarding$ forwarding-  
set policy policy-name set set-id <1..4>
```

An MPLS LSP can map only to a single class-forwarding policy and forwarding set. Multiple LSPs can map to the same policy and set. If they form an ECMP set, from the IGP shortcut perspective, packets of the FCs mapped to this set will be sprayed over these LSPs based on a modulo operation of the output of the hash routine on the packet's headers and the number of LSPs in the set.

2.3.2 Feature Behavior

When a BGP IPv4 or IPv6 prefix is resolved to a BGP next-hop, consisting of up to 64 resolved next-hops (LSPs and IP links), the default behavior of the data path is to spray the packets over the entire ECMP set using a modulo operation of the number of resolved next-hops in the ECMP set and the output of the hash on the packet header fields.

Both the CBF feature in LDP-over-RSVP and this CBF feature over IGP IPv4 shortcuts make use of the CBF class-forwarding policy. IGP always passes the CBF information populated by MPLS for each LSP used as a tunnel next-hop by an IGP prefix. The new CBF information is checked for consistency. If more than a single class-forwarding policy exists in the tunnel next-hops of a IGP prefix, IGP removes the new CBF information from all the corresponding tunnels and the behavior will be as if there were no CBF info.

When the CBF feature is enabled (**class-forwarding** option, enabled under **config>router** context), each application (BGP, CPM, or LDP), when looking up a prefix in RTM, will find up to 64 IP and tunnel next-hops. This lookup is split into three subsets:

- Subset 1 — tunnel next-hops with older CBF information (FCs mapped to this LSP, default LSP (true/false), CBF Policy ID=0, Set ID=any valid value). This information is usable by LDP only. Other applications treat this like non-CBF information.
- Subset 2 — tunnel next-hops with new CBF information (FCs mapped to this LSP, default LSP (true/false), CBF Policy ID>0, Set ID>0). This information is usable by both LDP and other applications.
- Subset 3 — tunnel-next-hops with no CBF information and IP next-hops. Usable by all applications, except that LDP will use tunnel next-hops only.

The BGP application performs a lookup in RTM for a prefix matching each BGP next-hop of a prefix. The BGP application selects tunnels belonging to the class-forwarding sets in Subset 2 and for each BGP next-hop of a prefix. The remaining tunnels, with no CBF configuration or with the older CBF information and the IP next-hops, are still programmed to IOM. However, BGP and the data path will use them only when all the class-forwarding sets are not available as explained below.

The SR OS implements a hierarchical ECMP architecture for BGP prefixes in the data path. The first level is the ECMP at the BGP next-hop level. The second level is ECMP at the resolved next-hop (IP or tunnel next-hop) level. The CBF feature is independently applied to the set of resolved tunnel next-hops of each BGP next-hop of a prefix. The user must make sure that the sets of LSPs that are used as IGP shortcuts to reach each of the BGP next-hops have the appropriate FC mappings.

The following procedures are enforced in the CBF feature.

- The tunnels in the full next-hop ECMP set, with set size greater or equal to 1 and less than or equal to 64, can use MPLS LSPs that terminate on multiple endpoints (BGP next-hop itself or otherwise) to reach the next-hop of a BGP prefix. The existing ECMP tunnel and IP next-hop selection behavior, when resolving a prefix over IGP shortcuts, continues to be used.

- If no LSP among the full ECMP set of a BGP next-hop has a class-forwarding policy configuration assigned, then the set is considered inconsistent from a CBF perspective. No CBF-related information is programmed in IOM and regular ECMP spraying over the full set occurs.
- If only a single class-forwarding policy is referenced by one or more LSPs in the full ECMP set of a BGP next-hop, the full set is considered consistent from a CBF perspective and the class-forwarding policy is used to spray packets of each FC over the LSPs within each forwarding set. As a result of this processing, only the LSPs that have been selected for forwarding traffic are programmed in IOM with CBF information. The remaining LSPs and IP next-hops of the BGP next-hop, are also programmed in IOM but without any CBF information associated and, therefore, will not be used for CBF.
- If multiple class-forwarding policies are referenced by LSPs in the full ECMP set of a BGP next-hop, the set is considered inconsistent from a CBF perspective. No CBF related information is programmed in IOM and regular ECMP spraying over the full set occurs.

The following describes the fallback behavior in data path of the CBF feature.

- An FC, for which all LSPs in the forwarding set are operationally down, has its packets forwarded over the default forwarding set. The default forwarding set is either the initial default forwarding set configured by the user or the lowest numbered set in the class-forwarding policy that has one or more LSPs in the operationally UP state. If the initial or subsequently elected default forwarding set has all its LSPs operationally down, the next lower numbered forwarding set, which has at least one LSP in the operationally up state, is elected as the default forwarding set.
- If all LSPs of all forwarding sets become operationally down, the router resumes regular ECMP spraying on the remaining LSPs and IP next-hops in the full ECMP set.
- Whenever the first LSP in a forwarding set becomes operationally UP, the router triggers the re-election of the default set and will select this set as the new default set, if it is the initial default set, otherwise, it will select lowest numbered set.

2.3.3 Feature Limitations

The following are the limitations of the CBF feature.

- CBF applies to packets of IPv4 and IPv6 BGP prefixes only. CBF does not apply to IGP prefixes and static route prefixes resolved over IGP IPv4 shortcuts. The latter are forwarded using regular ECMP over the entire set of up to 64 tunnel next-hops.
- CPM originated packets on the router, including control plane and OAM packets, are forwarded over a single LSP from the set of LSPs the packet's FC is mapped to, as per the CBF configuration. CPM, however, only maintains a maximum of 64 next-hops for a given destination prefix. Therefore, if there are multiple BGP next-hops for a prefix, CPM selects 64 tunnel next-hops by cycling over the BGP next-hops in ascending order. Then, the first LSP in the first set ID that the FC of the packet maps to is selected to forward the packet.

Furthermore, the CBF information consistency check, the CBF default set determination, and the CBF set failover procedures are applied to this set of 64 tunnel next-hops.

The user can configure the SGT-QoS feature to change the DSCP and FC of CPM-originated packets of a specific control plane protocol to select an LSP from a different set ID. This configuration allows, for instance, the forwarding of BGP Keep-Alive packets over an LSP of the same set ID as that of the data plane packets of the BGP prefixes destined to the same BGP next-hop.

- Weighted ECMP, at the transport tunnel level of BGP prefixes over IGP shortcuts, and the CBF feature on a per-BGP next-hop basis are mutually exclusive. Specifically, if the user enables both weighted ECMP (**configure>router>weighted-ecmp**) and CBF (**configure>router> class-forwarding**), weighted ECMP applies as long as all the LSPs used as tunnel next-hops to reach the BGP next-hop of a prefix have a user-configured weight. Otherwise, the CBF feature applies as per the procedures described in [Feature Behavior](#).

2.3.4 Data Path Support

When a packet of a BGP IPv4 or IPv6 prefix is received, the data path uses the FC that the packet was classified into to look up the forwarding set ID. The data path then performs a modulo operation on the tunnel next-hops of this set ID, to select the one next-hop for forwarding the packet. Therefore, packets matching an FC are only sprayed over the ECMP tunnel-next-hops of the set ID this FC maps to.

Both the BGP or CPM application and IOM use the same algorithm for failover and default class-forwarding set determination, as described in [Feature Behavior](#) and illustrated in [Example Configuration and Default CBF Set Election](#).

If MPLS deletes an LSP from a specified set ID, the IOM handles failover within the same set ID. The IOM reprograms the data path to spray packets of the impacted FCs over the remaining tunnel next-hops of the set ID.

Similarly, the IOM handles failover between class-forwarding sets when MPLS deletes the last LSP in a set ID. The IOM reprograms the data path to spray packets of the impacted FCs over the tunnel next-hops of the failover set ID. In both cases, the failover does not make use of the uniform failover procedure; however, if an LSP activated its FRR backup path, it remains in the set ID and continues to forward traffic of the mapped FCs.

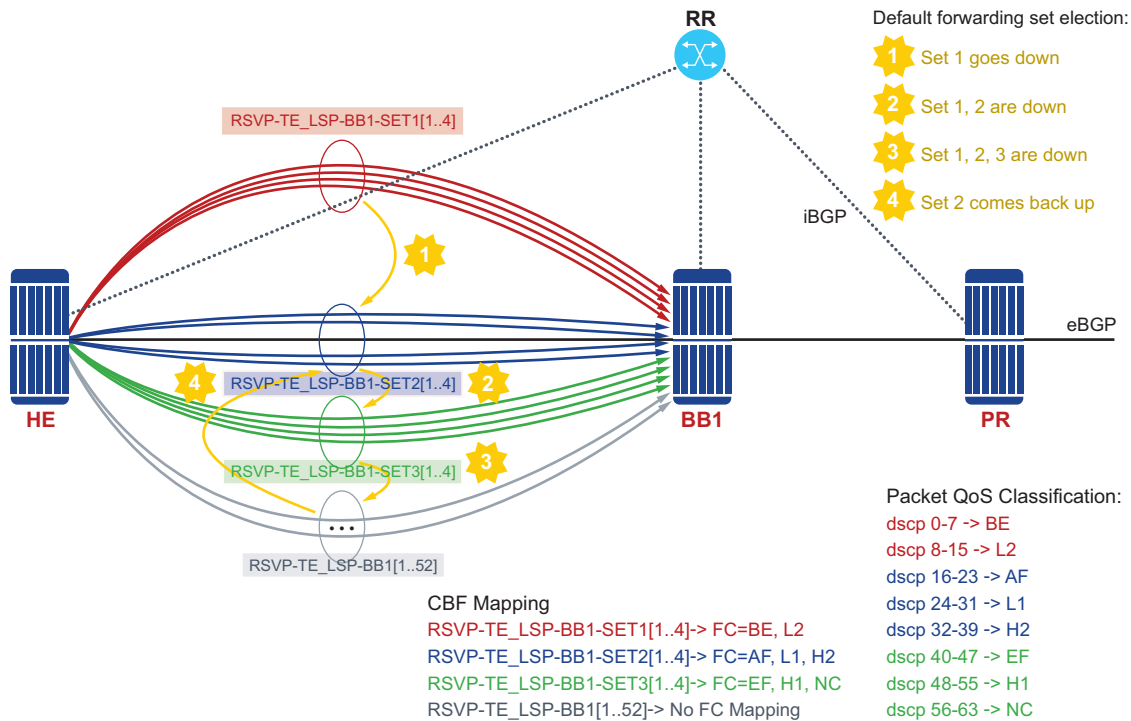
Finally, BGP updates the set IDs, used to reach a BGP next-hop, any time IGP updates the information in the RTM.

2.3.5 Example Configuration and Default CBF Set Election

Assume the following user configuration.

- The FC mapping to the sets and the default forwarding set election are illustrated in [Figure 10](#).
- All sets and RSVP-TE LSPs outside of the three class-forwarding sets are up initially.
- Set 1 is elected as the default class-forwarding set (because the user did not configure an initial default set).
- If All LSPs in Set 1 go operationally down, Set 2 is elected as the default class-forwarding set.
- If Set 2 subsequently goes down, Set 3 is elected as the default class-forwarding set.
- If Set 3 subsequently goes down, then packets of BGP prefixes will be ECMP sprayed over the remaining non-CBF RSVP-TE LSPs.
- If Set 2 comes back up, then Set 2 is elected as the default class-forwarding set.

Figure 10 Default Forwarding Set Election



sw0232

```
*A:Reno 194>config>router# info
-----
#-----
echo "IP Configuration"
#-----
interface "system"
  address 38.120.48.194/32
  ipv6
    address 3ffe::a14:194/128
  exit
  no shutdown
exit
interface "toSim199"
  address 10.202.5.194/24
  secondary 11.202.5.194/24
  port 1/1/1
  ipv6
    address 2001:db8:a0b:12f0::1/64
  exit
  no shutdown
exit
interface "toSim213"
  address 10.202.4.194/24
  port 1/1/2
  no shutdown
exit
interface "toSim219"
  address 10.202.8.194/24
  port 1/1/3
```

```

        no shutdown
    exit
    class-forwarding
    // Enables CBF feature for BGP and CPM traffic

*A:Reno 194>config>router>isis# info
-----
        igp-shortcut
//
    Enables IGP shortcut in this ISIS instance with both families IPv4 and IPv6 resolving to RSVP-TE LSPs
        tunnel-next-hop
            family ipv4
                resolution filter
                resolution-filter
                rsvp
            exit
        exit
        family ipv6
            resolution filter
            resolution-filter
            rsvp
        exit
    exit
    exit
    no shutdown
    exit
    no shutdown
-----
*A:Reno 194>config>router>mpls# info
-----
    class-forwarding-policy cbf1
        fc be forwarding-set 1
        fc l2 forwarding-set 1
        fc af forwarding-set 2
        fc l1 forwarding-set 2
        fc h2 forwarding-set 2
        fc ef forwarding-set 3
        fc h1 forwarding-set 3
        fc nc forwarding-set 3
    cspf-on-loose-hop
    exit
    interface "system"
        no shutdown
    exit
    interface "toSim199"
        no shutdown
    exit
    interface "toSim213"
        admin-group "olive"
        no shutdown
    exit
    interface "toSim219"
        no shutdown
    exit
    path "empty"
        no shutdown
    exit
    lsp "RSVP-TE_LSP-BB1-SET1[1..4]" // Four LSPs in Set1

```

```

        shutdown
        to 38.120.48.211
        cspf
        class-forwarding
            forwarding-set policy "cbf1" set 1
        exit
        primary "empty"
        exit
    exit
no shutdown
lsp "RSVP-TE_LSP-BB1-SET2[1..4]" // Four LSPs in Set2
    shutdown
    to 38.120.48.211
    cspf
    class-forwarding
        forwarding-set policy "cbf1" set 2
    exit
    primary "empty"
    exit
exit
lsp "RSVP-TE_LSP-BB1-SET3[1..4]" // Four LSPs in Set3
    shutdown
    to 38.120.48.211
    cspf
    class-forwarding
        forwarding-set policy "cbf1" set 3
    exit
    primary "empty"
    exit
exit
lsp "RSVP-TE_LSP-BB1[1..52]" //
Other LSP configuration with no CBF options for a total of 64 LSPs to BB1
    shutdown
    to 38.120.48.211
    cspf
    primary "empty"
    exit
exit
no shutdown
-----

```

2.4 Bidirectional Forwarding Detection

Bidirectional Forwarding Detection (BFD) is an efficient, short-duration detection of failures in the path between two systems. If a system stops receiving BFD messages for a long enough period (based on configuration), it is assumed that a failure along the path has occurred and the associated protocol or service is notified of the failure.

BFD can provide a mechanism used for failure detection over any media, at any protocol layer, with a wide range of detection times and overhead, to avoid a proliferation of different methods.

SR OS supports asynchronous and on-demand modes of BFD in which BFD messages are sent to test the path between systems.

If multiple protocols are running between the same two BFD endpoints, only a single BFD session is established, and all associated protocols will share the single BFD session.

As well as the typical asynchronous mode, there is also an echo function defined within RFC 5880, *Bidirectional Forwarding Detection*, that allows either of the two systems to send a sequence of BFD echo packets to the other system, which loops them back within that system's forwarding plane. If a number of these echo packets are lost, the BFD session is declared down.

2.4.1 BFD Control Packet

The base BFD specification does not specify the encapsulation type to be used for sending BFD control packets. Instead, use the appropriate encapsulation type for the medium and network. The encapsulation for BFD over IPv4 and IPv6 networks is specified in draft-ietf-bfd-v4v6-1hop-04.txt, *BFD for IPv4 and IPv6 (Single Hop)*. This specification requires that BFD control packets be sent over UDP with a destination port number of 3784 and the source port number must be within the range 49152 to 65535.

Also, the TTL of all transmitted BFD packets must have an IP TTL of 255. All BFD packets received must have an IP TTL of 255 if authentication is not enabled. If authentication is enabled, the IP TTL should be 255, but can still be processed if it is not (assuming the packet passes the enabled authentication mechanism).

If multiple BFD sessions exist between two nodes, the BFD discriminator is used to de-multiplex the BFD control packet to the appropriate BFD session.

2.4.2 Control Packet Format

The BFD control packet has two sections: a mandatory section and an optional authentication section.

Figure 11 Mandatory Frame Format

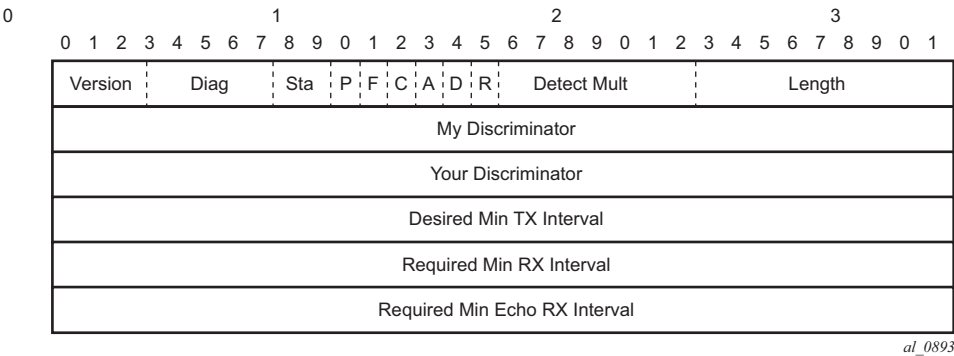


Table 5 BFD Control Packet Field Descriptions

| Field | Description |
|-------|---|
| Vers | The version number of the protocol. The initial protocol version is 0. |
| Diag | A diagnostic code specifying the local system's reason for the last transition of the session from Up to some other state. Possible values are: 0-No diagnostic 1-Control detection time expired 2-Echo function failed 3-Neighbor signaled session down 4-Forwarding plane reset 5-Path down 6-Concatenated path down 7-Administratively down |
| D Bit | The demand mode bit. (Not supported) |
| P Bit | The poll bit. If set, the transmitting system is requesting verification of connectivity, or of a parameter change. |
| F Bit | The final bit. If set, the transmitting system is responding to a received BFD control packet that had the poll (P) bit set. |
| Rsvd | Reserved bits. These bits must be zero on transmit and ignored on receipt. |

Table 5 BFD Control Packet Field Descriptions (Continued)

| Field | Description |
|-------------------------------|--|
| Length | Length of the BFD control packet, in bytes. |
| My Discriminator | A unique, non-zero discriminator value generated by the transmitting system, used to demultiplex multiple BFD sessions between the same pair of systems. |
| Your Discriminator | The discriminator received from the corresponding remote system. This field reflects back the received value of my discriminator, or is zero if that value is unknown. |
| Desired Min TX Interval | This is the minimum interval, in microseconds, that the local system would like to use when transmitting BFD control packets. |
| Required Min RX Interval | This is the minimum interval, in microseconds, between received BFD control packets that this system is capable of supporting. |
| Required Min Echo RX Interval | This is the minimum interval, in microseconds, between received BFD echo packets that this system is capable of supporting. If this value is zero, the transmitting system does not support the receipt of BFD echo packets. |

2.4.3 BFD for RSVP-TE

BFD will notify RSVP-TE if the BFD session goes down, in addition to notifying other configured BFD enabled protocols (for example, OSPF, IS-IS, and PIM). This notification will then be used by RSVP-TE to begin the reconvergence process. This greatly accelerates the overall RSVP-TE response to network failures.

All encapsulation types supporting IPv4 and IPv6 are supported because all BFD packets are carried in IPv4 and IPv6 packets; this includes Frame Relay and ATM.

BFD is supported on the following interfaces:

- Ethernet (Null, Dot1Q & QinQ)
- Spoke SDPs
- LAG interfaces

The following interfaces are supported only on the 7750 SR and 7450 ESS:

- VSM interfaces
- POS interfaces (including APS)
- Channelized interfaces (PPP, HDLC, FR, and ATM) on ASAP (priority 1) and channelized MDAs (priority 2) including link bundles and IMA

2.4.4 Echo Support

Echo support for BFD calls for the support of the echo function within BFD. By supporting BFD echo, the router loops back received BFD echo messages to the original sender based on the destination IP address in the packet.

The echo function is useful when the local router does not have sufficient CPU power to handle a periodic polling rate at a high frequency. Therefore, it relies on the echo sender to send a high rate of BFD echo messages through the receiver node, which is only processed by the receiver's forwarding path. This allows the echo sender to send BFD echo packets at any rate.

SR OS does not support the sending of echo requests, only the response to echo requests.

2.4.5 BFD Support for BGP

This feature allows BGP peers to be associated with the BFD session. If the BFD session fails, BGP peering will also be torn down.

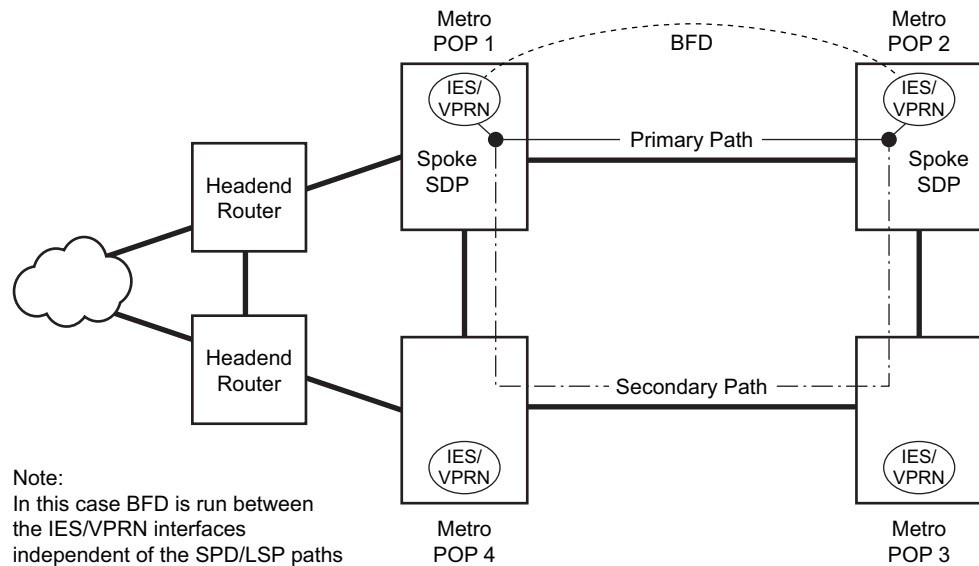
2.4.6 Centralized BFD

The following applications of centralized BFD require BFD to run on the SF/CPM.

2.4.6.1 IES Over Spoke SDP

One application for a central BFD implementation is so BFD can be supported over spoke SDPs used to inter-connect IES or VPRN interfaces. When there are spoke SDPs for inter-connections over an MPLS network between two routers, BFD is used to speed up failure detections between nodes so re-convergence of unicast and multicast routing information can begin as quickly as possible.

The MPLS LSP associated with the spoke SDP can enter or egress from multiple interfaces on the router. BFD for these types of interfaces cannot exist on the IOMXCM by itself.


Figure 12 BFD for IES/VPRN over Spoke SDP

Fig_31

2.4.6.2 BFD Over LAG and VSM Interfaces

A second application for a central BFD implementation is so BFD can be supported over LAG or VSM interface. This is useful where BFD is not used for link failure detection, but for node failure detection. In this application, the BFD session can run between the IP interfaces associated with the LAG or VSM interface, but there is only one session between the two nodes. There is no requirement for the message flow to across a certain link, or VSM, to get to the remote node.

Figure 13 BFD Over LAG and VSM Interfaces



Fig_32

2.4.6.3 LSP BFD and VCCV BFD

BFD is supported over MPLS-TP, RSVP, and LDP LSPs, as well as over pseudowires that support Layer 2 services such as Epipe VPLS spoke-SDPs and mesh-SDPs using centralized BFD. See the *7450 ESS, 7750 SR, 7950 XRS, and VSR MPLS Guide* and *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for more information.

2.4.7 Aggregate Next Hop

This feature adds the ability to configure an indirect next-hop for aggregate routes. The indirect next-hop specifies where packets will be forwarded if they match the aggregate route, but is not a more-specific route in the IP forwarding table.

2.4.8 Invalidate Next-Hop Based on ARP/Neighbor Cache State

This feature invalidates next-hop entries for static routes when the next-hop is no longer reachable on directly connected interfaces. This invalidation is based on ARP and Neighbor Cache state information.

When a next-hop is detected as no longer reachable due to ARP/Neighbor Cache expiry, the route's next-hop is set as unreachable to prevent the SR from sending continuous ARPs/Neighbor Solicitations triggered by traffic destined for the static route prefix. When the next-hop is detected as reachable via ARP or Neighbor Advertisements, the state of the next-hop is set back to valid.

2.4.8.1 Invalidate Next-Hop Based on IPV4 ARP

This feature invalidates a static route based on the reachability of the next-hop in the ARP cache when the **validate-next-hop** command is enabled within the **static-route-entry>next-hop** context for an IPv4 static route.

In this case, when the ARP entry for the next-hop is INVALID or not populated, the static route must remain invalid/inactive. When an ARP entry for the next-hop is populated based on a gratuitous ARP received or periodic traffic destined for it and the usual ARP who-has procedure, the static route becomes valid/active and is installed.

2.4.8.2 Invalidate Next-Hop Based on Neighbor Cache State

This feature invalidates a static route based on the reachability of the next-hop in the neighbor cache when the **validate-next-hop** command is enabled within the **static-route-entry>next-hop** context for an IPv6 static route.

In this case, when the Neighbor Cache entry for next-hop is INVALID or not populated, the static route must remain invalid/inactive. When an NC entry for next-hop is populated based on a neighbor advertisement received, or periodic traffic destined for it and the usual NS/NA procedure, the static route becomes valid/active and is installed.

2.4.9 LDP Shortcut for IGP Route Resolution

This feature enables you to forward user IP packets and specified control IP packets using LDP shortcuts over all network interfaces in the system that participate in the IS-IS and OSPF routing protocols. The default is to disable the LDP shortcut across all interfaces in the system.

```
config>router>ldp-shortcut [ipv4] [ipv6]
```

2.4.9.1 IGP Route Resolution

When LDP shortcut is enabled, LDP populates the RTM with next-hop entries corresponding to all prefixes for which it activated an LDP FEC. For an activated prefix, two route entries are populated in RTM. One corresponds to the LDP shortcut next-hop and has an owner of LDP. The other one is the regular IP next-hop. The LDP shortcut next-hop always has preference over the regular IP next-hop for forwarding user packets and specified control packets over a specific outgoing interface to the route next-hop.

The prior activation of the FEC by LDP is done by performing an exact match with an IGP route prefix in RTM. It can also be done by performing a longest prefix match with an IGP route in RTM if the aggregate-prefix-match option is enabled globally in LDP *ldp-interarea-prd*.

The LDP next-hop entry is not exported to the LDP control plane or to any other control plane protocols except OSPF, IS-IS, and an OAM control plane specified in [Handling of Control Packets](#).

This feature is not restricted to /32 IPv4 prefixes or /128 IPv6 FEC prefixes. However, only /32 IPv4 and /128 IPv6 FEC prefixes will be populated in the tunnel table for use as a tunnel by services.

All user and specified control packets for which the longest prefix match in RTM yields the FEC prefix will be forwarded over the LDP LSP. The following is an example of the resolution process.

Assume that the egress LER advertised a FEC for some /24 prefix using the fec-originate command. At the ingress LER, LDP resolves the FEC by checking in RTM that an exact match exists for this prefix. After the LDP activates the FEC, it programs the NHLFE in the egress data path and the LDP tunnel information in the ingress data path tunnel table.

Next, LDP provides the shortcut route to RTM, which will associate it with the same /24 prefix. There will be two entries for this /24 prefix: the LDP shortcut next-hop and the regular IP next-hop. The latter was used by LDP to validate and activate the FEC. RTM then resolves all user prefixes that succeed a longest prefix match against the /24 route entry to use the LDP LSP.

Now assume that the aggregate-prefix-match was enabled and that LDP found a /16 prefix in RTM to activate the FEC for the /24 FEC prefix. In this case, RTM adds a new, more-specific route entry of /24 and has the next-hop as the LDP LSP. However, RTM will still not have a specific /24 IP route entry. RTM then resolves all user prefixes that succeed a longest prefix match against the /24 route entry to use the LDP LSP. All other prefixes that succeed a longest prefix match against the /16 route entry will use the IP next-hop. LDP shortcut will also work when using RIP for routing.

2.4.9.2 LDP-IGP Synchronization

See the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR MPLS Guide* for information about LDP-IGP Synchronization.

2.4.9.3 LDP Shortcut Forwarding Plane

After the LDP activates a FEC for a prefix and programs RTM, it also programs the ingress tunnel table in IOM or on linecards with the LDP tunnel information.

When an IPv4 packet is received on an ingress network interface, a subscriber IES interface, or a regular IES interface, the lookup of the packet by the ingress IOM or linecard will result in the packet being sent labeled with the label stack corresponding to the NHLFE of the LDP LSP when the preferred RTM entry corresponds to an LDP shortcut.

If the preferred RTM entry corresponds to an IP next-hop, the IPv4 packet is forwarded unlabeled.

The switching from the LDP shortcut next-hop to the regular IP next-hop when the LDP FEC becomes unavailable depends on whether the next-hop is still available. If it is (for example, the LDP FEC was withdrawn due to LDP control plane issues) the switchover should be faster. If the next-hop determination requires IGP to re-converge, this will take longer. However, no target is set.

The switching from a regular IP next-hop to an LDP shortcut next-hop will usually occur only when both are available. However, the programming of the NHLFE by LDP and the programming of the LDP tunnel information in the ingress IOM or linecards tunnel table are asynchronous. If the tunnel table is configured first, it is possible that traffic will be black-holed for some time.

2.4.9.4 ECMP Considerations

When ECMP is enabled and multiple equal-cost next-hops exist for the IGP route, the ingress IOM or linecard will spray the packets for this route based on the hashing routine currently supported for IPv4 packets.

When the preferred RTM entry corresponds to an LDP shortcut route, spraying will be performed across the multiple next-hops for the LDP FEC. The FEC next-hops can either be direct link LDP neighbors or T-LDP neighbors reachable over RSVP LSPs, in the case of LDP-over-RSVP, but not both. This is as per ECMP for LDP.

When the preferred RTM entry corresponds to a regular IP route, spraying will be performed across regular IP next-hops for the prefix.

Spraying across regular IP next-hops and LDP-shortcut next-hops concurrently is not supported.

2.4.9.5 Handling of Control Packets

All control plane packets will not see the LDP shortcut route entry in RTM with the exception of the following control packets, which will be forwarded over an LDP shortcut when enabled:

- A locally generated or in transit ICMP ping and trace route of an IGP route. The transit message appears as a user packet to the ingress LER node.
- A locally generated response to a received ICMP ping or trace route message.

All other control plane packets that require an RTM lookup and knowledge of which destination is reachable over the LDP shortcut will continue to be forwarded over the IP next-hop route in RTM.

2.4.9.6 Handling of Multicast Packets

Multicast packets cannot be forwarded or received from an LDP LSP. This is because there is no support for the configuration of such an LSP as a tunnel interfaces in PIM. Only an RSVP P2MP LSP is currently allowed.

If a multicast packet is received over the physical interface, the uRPF check will not resolve to the LDP shortcut because the LDP shortcut route in RTM is not made available to multicast application.

2.4.9.7 Interaction with BGP Route Resolution to an LDP FEC

There is no interaction between an LDP shortcut for BGP next-hop resolution and the LDP shortcut for IGP route resolution. BGP will continue to resolve a BGP next-hop to an LDP shortcut if the user enabled the following option in BGP:

```
config>router>bgp>next-hop-res>shortcut-tunnel
      family ipv4
      resolution-filter ldp
```

2.4.9.8 Interaction with Static Route Resolution to an LDP FEC

A static route will continue to be resolved by searching an LDP LSP whose FEC prefix matches the specified indirect next-hop for the route. In contrast, the LDP shortcut for IGP route resolution uses the LDP LSP as a route. The most specific route for a prefix will be selected and, if both a static and IGP routes exist, the RTM route type preference will be used to select one.

2.4.9.9 LDP Control Plane

For the LDP shortcut to be usable, SR OS must originate a <FEC, label> binding for each IGP route it learns of even if it did not receive a binding from the next-hop for that route. The router must assume that it is an egress LER for the FEC until the route disappears from the routing table or the next-hop advertises a binding for the FEC prefix. In the latter case, SR OS becomes a transit LSR for the FEC.

SR OS will originate a <FEC, label> binding for its system interface address only by default. The only way to originate a binding for local interfaces and routes that are not local to the system is by using the fec-originate capability.

You must use the **fec-originate** command to generate bindings for all non-local routes for which this node acts as an egress LER for the corresponding LDP FEC. Specifically, this feature must support the FEC origination of IGP learned routes and subscriber/host routes statically configured or dynamically learned over subscriber IES interfaces.

An LDP LSP used as a shortcut by IPv4 packets may also be tunneled using the LDP-over-RSVP feature.

2.5 Weighted Load-Balancing over Interface Next-hops

When the **weighted-ecmp** command is configured in the base router context (**config>router**) or a VPRN (**config>service>vprn**), the associated IS-IS instances are allowed to program IPv4 and IPv6 ECMP routes to use weighted load-balancing across interface next-hops. The following conditions must be true:

- All ECMP next-hops must be interface next-hops.
- All ECMP next-hops must be associated with the same neighbor IS-IS router.
- All ECMP next-hop interfaces must have a non-zero **load-balancing-weight** value configured in the **config>router>isis>interface** context.

By default, IS-IS interfaces have a zero weight (**no load-balancing-weight**); non-zero values must be configured explicitly. Values cannot be auto-derived.

The **config>router>isis>interface>load-balancing-weight** command accepts a value between 0 and 4294967295.

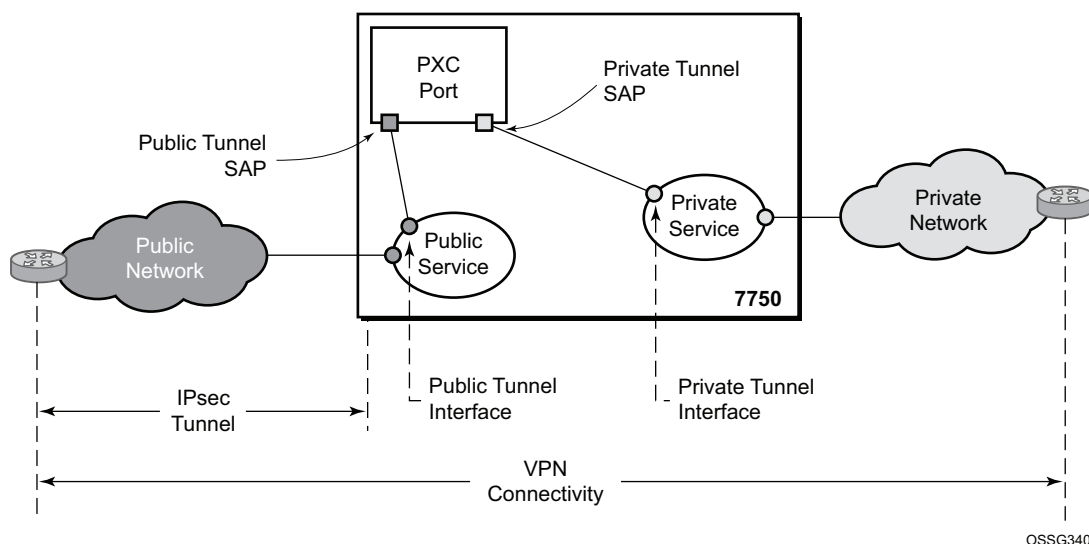
If a base router IPv4 or IPv6 BGP route has a BGP next-hop resolved by an ECMP IS-IS route and **ibgp-multipath** is configured under BGP, traffic forwarded to the BGP next-hop is sprayed according to the load-balancing weights of the interface next-hops.

2.6 GRE Tunnel Overview

This section describes the GRE tunneling feature supported through the use of a Port Cross Connect (PXC) port. In this application, the PXC port functions as a resource module for the system, providing the necessary resources for the GRE encapsulation function. The GRE encapsulation function described here is similar to the GRE tunnel functionality supported through the use of the MS-ISA. In this use case, the MS-ISA is not required.

Figure 14 shows an example of a GRE deployment supported inside a 7750 SR router using the PXC element.

Figure 14 Sample GRE Deployment Using a PXC Port



In Figure 14, the public network is typically an unsecured network, such as public Internet, over which packets belonging to the private network in the diagram cannot be transmitted natively. Inside the 7750 SR, a public service instance (IES or VRPN) connects to the public network, and a private service instance (typically a VRPN) connects to the private network.

For GRE tunnels using PXC ports, the public and private services must be two different services, and the PXC is the connection between the two services. Traffic from the public network may require authentication and encryption inside an IPsec tunnel to reach the private network. In this way, the authenticity, confidentiality, and integrity of private network access can be enforced. If authentication and confidentiality are not required, then access to the private network may be provided through GRE or IP-IP tunnels.

Traffic flows through PXC-based tunnels in the following ways:

- In the upstream direction (public to private), the encapsulated traffic is forwarded to a public tunnel interface if the destination address matches the local or gateway address of a GRE tunnel. As the traffic passes through the PXC port, the tunnel header is removed, the payload IP packet is delivered to the private service, and from there, the traffic is forwarded again based on the destination address of the payload IP packet.
- In the downstream direction (private to public), unencapsulated traffic belonging to the private service is forwarded into the tunnel by matching a route with the GRE tunnel as next-hop. The route can be configured statically, learned by running OSPF on the private tunnel interface or by running BGP over the tunnel. After clear traffic is forwarded to the PXC port, it is encapsulated in the GRE header and passed to the public service, and from there, the traffic is forwarded again based on the destination address of the GRE header.

2.6.1 Sample GRE Tunnel Configurations

Public interface example:

```
config > service > ies 100
  interface "int-gre-tunnel-public" create
    address 192.110.1.1/30
    sap pxc-1.b:100 create //Public interface
    description "Public Tunnel PXC SAP"
    exit
  exit
no shutdown
```

Private interface example:

```
config > service > vprn 200 customer 200 create
  route-distinguisher 64496:1
  vrf-target target:64496:1
  interface "int-gre-tunnel-private" tunnel create // Private if
    address 10.1.1.1/30
    ip-mtu 1476
    sap pxc-1.a:200 create
      ip-tunnel "gre-tunnel-1" create
        source 192.110.1.2
        remote-ip 192.120.1.1
        backup-remote-ip 192.120.1.2
        delivery-service 100
        gre-header send-key 123 receive-key 123
        no shutdown
      exit
    exit
  static-route 172.16.1.1/24 next-hop 10.1.1.2
... [additional SAPs and or SDP configuration]
```

2.7 Router Interface Encryption with NGE


NGE nodes support Layer 3 encryption on router interfaces for IPv4 traffic. NGE is not supported on dual-stack IPv4/IPv6 or IPv6-only interfaces. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Services Overview Guide* for more information about platforms that support NGE.

NGE is enabled on a router interface by configuring the **group-encryption** command on the router interface. The interface is considered part of the NGE domain, and any received packets that are NGE-encrypted are decrypted if the key group is configured on the node. To encrypt packets egressing the interface, the outbound key group must be configured on the interface. All IP packets, such as self-generated traffic or packets forwarded from router interfaces that are not inside the NGE domain, are encrypted when egressing the interface. There are some exceptions to this general behavior, as described in the sections below; for example, GRE-MPLS and MPLSoUDP packets are not encrypted when router interface encryption is enabled.

The outbound and inbound key groups configured on the router interface determine which keys are used to encrypt and decrypt traffic. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Services Overview Guide* for more information about configuring key groups.

To perform encryption, router interface encryption reuses the IPsec transport mode packet format as shown in [Figure 15](#).

Figure 15 Router Interface Encryption Packet Format (IPsec Transport Mode)



The protocol field in the IP header of an NGE packet is always set to “ESP”. Within an NGE domain, the SPI that is included in the ESP header is always an SPI for the key group configured on the router interface. Other fields in the IP header, such as the source and destination addresses, are not altered by NGE router interface encryption. Packets are routed through the NGE domain and decrypted when the packet leaves the NGE domain.

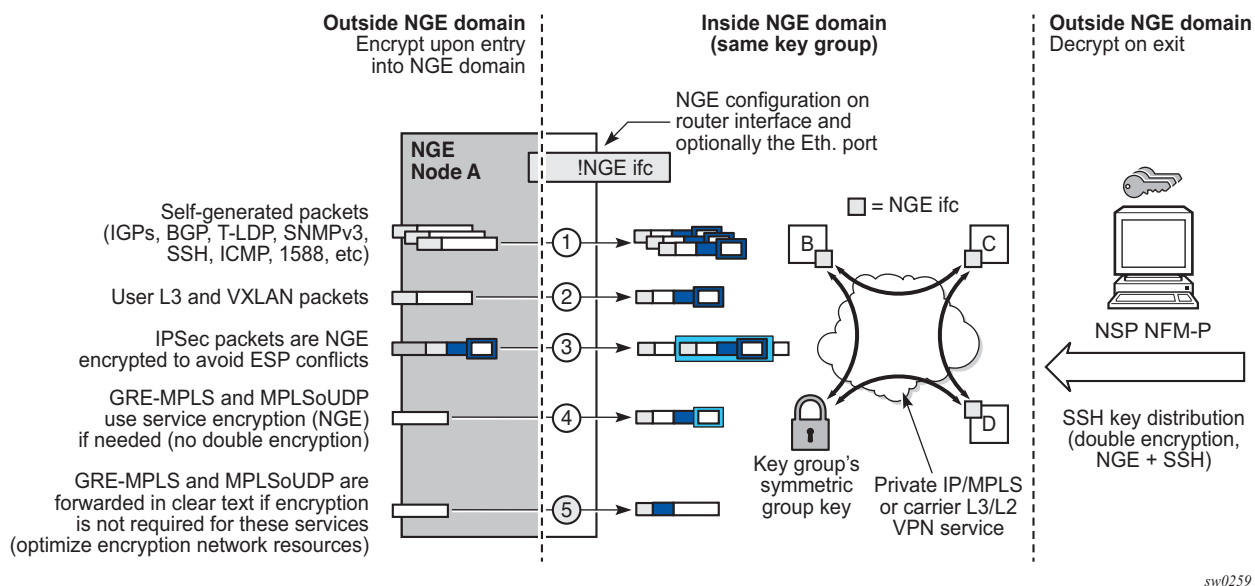
The group keys used on an NGE-enabled router interface provide encryption of broadcast and multicast packets within the GRT. For example, OSPF uses a broadcast address to establish adjacencies, which can be encrypted by NGE without the need to establish point-to-point encryption tunnels. Similarly, multicast packets are also encrypted without point-to-point encryption tunnels.

2.7.1 NGE Domains

An NGE domain is a group of nodes and router interfaces forming a network that uses a single key group to create a security domain. NGE domains are created when router interface encryption is enabled on router interfaces that need to participate in the NGE domain. The NSP NFM-P assists operators in managing the nodes and interfaces that participate in the NGE domain. See the *NSP NFM-P User Guide* for more information.

Figure 16 shows various traffic types crossing an NGE domain.

Figure 16 NGE Domain Transit



In Figure 16, nodes A, B, C, and D have router interfaces configured with router interface encryption enabled. Traffic is encrypted when entering the NGE domain using the key group configured on the router interface and is decrypted when exiting the NGE domain. Traffic may traverse multiple hops before exiting the NGE domain, yet decryption only occurs on the final node when the traffic exits the NGE domain.

Various traffic types are supported and encrypted when entering the NGE domain, as illustrated by the following items on node A in [Figure 16](#):

- item 1: self-generated packets — these packets, which include all types of control plane and management packets such as OSPF, BGP, LDP, SNMPv3, SSH, ICMP, RSVP-TE, and 1588, are encrypted
- item 2: user Layer 3 and VXLAN packets — any Layer 3 user packets that are routed into the NGE domain from an interface outside the NGE domain are encrypted. Any VXLAN packets that are routed into the NGE domain from this NGE node are encrypted.
- item 3: IPSec packets — IPSec packets are NGE-encrypted when entering the NGE domain to ensure that the IPSec packets' security association information does not conflict with the NGE domain

GRE-MPLS- or MPLSoUDP-based service traffic consists of Layer 3 packets, and router interface NGE is not applied to these types of packets. Instead, service-level NGE is used for encryption to avoid double-encrypting these packets and impacting throughput and latencies. The two types of GRE-MPLS or MPLSoUDP packets that can enter the NGE domain are illustrated by items 4 and 5 in [Figure 16](#).


- item 4: GRE-MPLS and MPLSoUDP packets (SDP or VPRN) with service-level NGE enabled — these encrypted packets use the key group that is configured on the service. The services key group may be different from the key group configured on the router interface where the GRE-MPLS or MPLSoUDP packet enters the NGE domain.
- item 5: GRE-MPLS and MPLSoUDP packets (SDP or VPRN) with NGE disabled — these packets are not encrypted and can traverse the NGE domain in clear text. If these packets require encryption, SDP or VPRN encryption must be enabled.

Creating an NGE domain from the NSP NFM-P requires the operator to determine the type of NGE domain being managed. This will indicate whether NGE gateway nodes are required to manage the NGE domain, and other operational considerations. The two types of NGE domains are:

- [Private IP/MPLS Network NGE Domain](#)
- [Private Over Intermediary Network NGE Domain](#)

2.7.1.1 Private IP/MPLS Network NGE Domain

One type of NGE domain is a private IP/MPLS network, as shown in [Figure 17](#).

Figure 17 Private IP/MPLS Network NGE Domain

26215


In a private IP/MPLS network NGE domain, all interfaces are owned by the operator and there is no intermediary service provider needed to interconnect nodes. Each interface is a point-to-point private link between private nodes. When a new node is added to this type of NGE domain (node D in [Figure 17](#)), the links that connect node D to the existing nodes in the NGE domain (nodes A, B, and C) must be enabled with NGE router interface encryption. Links from the new node to the existing nodes are enabled one at a time. The NSP NFM-P provides tools that simplify adding nodes to the NGE domain and enabling NGE on their associated interfaces. In this type of NGE domain, each interface is a direct link between two nodes and is not used to communicate with multiple nodes over a broadcast medium offered by an intermediary network. Also, there are no NGE gateway nodes required between the NSP NFM-P and new nodes entering the NGE domain.

2.7.1.2 Private Over Intermediary Network NGE Domain

The other type of NGE domain is a private IP/MPLS network that traverses an intermediary network NGE domain; the intermediary network is used to interconnect nodes in the NGE domain using a multipoint-to-multipoint service. The intermediary network is typically a service provider network that provides a private IP VPN service or a private VPLS service used to interconnect a private network that does not mimic point-to-point links as described in the [Private IP/MPLS Network NGE Domain](#) section.

This type of NGE domain is shown in [Figure 18](#).

Figure 18 Private Over Intermediary Network NGE Domain



26214

Private over intermediary network NGE domains have nodes with links that connect to a service provider network where a single link can communicate with multiple nodes over a Layer 3 service such as a VPRN. In [Figure 18](#), node A has NGE enabled on its interface with the service provider and uses that single interface to communicate with nodes B and C, and eventually with node D when node D has been added to the NGE domain. This type of NGE domain requires the recognition of NGE gateway nodes that allow the NSP NFM-P to reach new nodes that enter the domain. Node C is designated as a gateway node.

When node D is added to the NGE domain, it must first have the NGE domain key group downloaded to it from the NSP NFM-P. The NSP NFM-P creates an NGE exception ACL on the gateway node, C, to allow communication with node D using SNMPv3 and SSH through the NGE domain. After the key group is downloaded, the NSP NFM-P enables router interface encryption on node D's interface with the service provider and node D is now able to participate in the NGE domain. The NSP NFM-P automatically removes the IP exception ACL from node C when node D enters the NGE domain.

See [Router Interface NGE Domain Concepts](#) for more information.

2.7.2 Router Interface NGE Domain Concepts

An NGE domain is a group of nodes whose router interfaces in the base routing context (GRT) are enabled for router interface NGE. An interface without router interface NGE enabled is considered to be outside the NGE domain. NGE domains use only one key group when the domain is created; however, two key groups may be active at once if some links within the NGE domain are in transition from one key group to the other.

Figure 19 illustrates the NGE domain concept. Table 6 describes the three configuration scenarios inside the NGE domain.

Figure 19 Inside and Outside NGE Domains

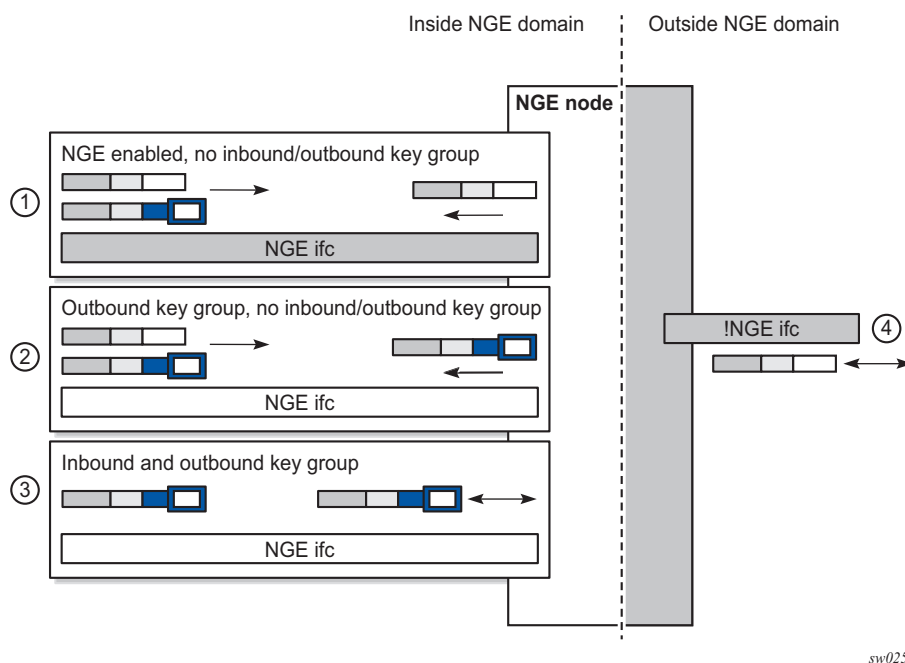


Table 6 Inside and Outside NGE Domains — Configuration Scenarios

| Key | Description |
|-----|--|
| 1 | NGE enabled, no inbound/outbound key group Outbound packets are sent without encrypting. Inbound packets can be NGE-encrypted or clear text. |
| 2 | Outbound key group, no inbound key group Outbound packets are encrypted using the interface key group if not already encrypted. Inbound packets can be NGE-encrypted or clear text. |

Table 6 Inside and Outside NGE Domains — Configuration Scenarios

| Key | Description |
|-----|--|
| 3 | Inbound and outbound key group Outbound packets are encrypted using the interface key group if not already encrypted. Inbound packets must be encrypted using the interface key group keys. |
| 4 | Outside the NGE domain, the interface is not configured for NGE. Any ESP packets are IPSec packets. |

A router interface is considered to be inside the NGE domain when it has been configured with **group-encryption** on the interface. When **group-encryption** is configured on the interface, the router can receive unencrypted packets or NGE-encrypted packets from any configured key group on the router, but any other type of IPSec-formatted packet is not allowed. If an IPSec-formatted packet is received on an interface that has **group-encryption** enabled, it will not pass NGE authentication and will be dropped. Therefore, IPSec packets cannot exist within the NGE domain without first being converted to NGE packets. This conversion requirement delineates the boundary of the NGE domain and other IPSec services.

When NGE router interface encryption is enabled and only an outbound key group is configured, the interface can receive unencrypted packets or NGE-encrypted packets from any configured key group on the router. All outbound packets are encrypted using the outbound key group if the packet was not already encrypted further upstream in the network.

When NGE router interface encryption has been configured with both an inbound and outbound key group, only NGE packets encrypted with the key group security association can be sent and received over the interface.

When there is no NGE router interface encryption, the interface is considered outside the NGE domain where NGE is not applied.

Refer to the “NGE Packet Overhead and MTU Considerations” section in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Services Overview Guide* for MTU information related to enabling NGE on a router interface.

2.7.3 GRE-MPLS and MPLSoUDP Packets Inside the NGE Domain

NGE router interface encryption is never applied to GRE-MPLS or MPLSoUDP packets, for example:

- GRE with the GRE protocol ID set to MPLS Unicast (0x8847) or Multicast (0x8848)
- UDP packets with destination port = 6635)

GRE-MPLS and MPLSoUDP packets that enter the NGE domain or transit the NGE domain are forwarded as is.

Because these GRE-MPLS and MPLS-oUDP packets provide transport for MPLS-based services, they already use the NGE services-based encryption techniques for MPLS, such as SDP or VPRN-based encryption. To avoid double encryption, the packets are left in clear text when entering an NGE domain or crossing intermediate nodes in the NGE domain, and are forwarded as needed when exiting an NGE domain.

2.7.4 EVPN-VXLAN Tunnels and Services

NGE router interface encryption does not differentiate between EVPN-VXLAN tunnels and other L3 traffic, and therefore encrypts all EVPN-VXLAN traffic that egresses the node.

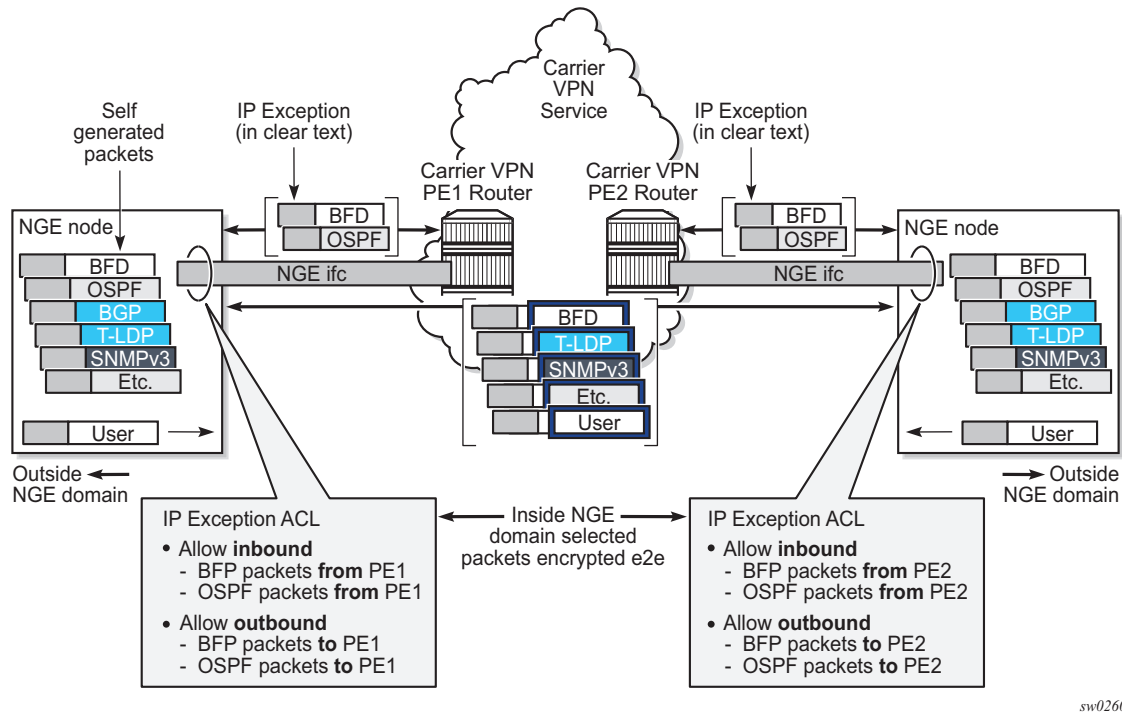
For received encrypted EVPN-VXLAN packets, if the VXLAN tunnel terminates on the node (that is, the destination IP is for a VTEP on this node), then the NGE packet is decrypted and the EVPN-VXLAN traffic is processed as if NGE encryption never took place.

2.7.5 Router Encryption Exceptions using ACLs

In some cases, Layer 3 packets may need to cross the NGE domain in clear text, such as when an NGE-enabled router needs to peer with a non-NGE-capable router to exchange routing information. This can be accomplished by using a router interface NGE exception filter applied on the router interface for the required direction, inbound or outbound.

[Figure 20](#) shows the use of a router interface NGE exception filter.

Figure 20 Router Interface NGE Exception Filter Example



The inbound or outbound exception filter is used to allow specific packet flows through the NGE domain in clear text, where there is an explicit inbound and outbound key group configured on the interface. The behavior of the exception filter for each router interface configuration is as follows:

- NGE enabled, no inbound/outbound key group — in this scenario, the router does not encrypt outbound traffic, and so the outbound exception filter is not applied. The router can still receive inbound NGE packets, so the exception filter is applied to inbound packets. If the filter detects a match, clear text packets can be received and forwarded by the router.
- outbound key group, no inbound key group — the outbound exception filter is applied to outbound traffic, and packets that match the filter are not encrypted on egress. The router can receive inbound NGE packets without an inbound key group set and applies the exception filter to inbound packets. If the filter detects a match, clear text packets can be received and forwarded by the router.
- inbound and outbound key group — the inbound and outbound exception filters are applied, and any packets that match are passed in clear text.


2.7.6 IPSec Packets Crossing an NGE Domain

IPSec packets can cross the NGE domain because they are still considered Layer 3 packets. To avoid confusion between the security association used in an IPSec packet and the one used in a router interface NGE packet, the router will always apply NGE to any IPSec packet that traverses the NGE domain.

IPSec packets that originate from a router within the NGE domain are not allowed to enter the NGE domain. The only exception to this restriction is OSPFv3 packets.

Figure 21 shows how IPSec packets can transit an NGE domain.

Figure 21 IPSec Packets Transiting an NGE Domain



sw0256

An IPSec packet enters the router from outside the NGE domain. When the router determines that the egress interface to route the packet is inside an NGE domain, it will select an NGE router interface with one of the following configurations.

- NGE enabled with no inbound or outbound key group configured — this link cannot forward the IPSec packet without adding the NGE ESP, but since nothing is configured for the outbound key group, the packet must be dropped.

- NGE enabled with outbound key group configured and no inbound key group configured — the packet originates outside the NGE domain, so the router adds an ESP header over the existing ESP and encrypts the payload using the NGE domain keys for the configured outbound key group.
- NGE enabled with both inbound and outbound key groups configured — the packet originates outside the NGE domain, so the router adds an ESP header over the existing ESP and encrypts the payload using the NGE domain keys for the configured outbound key group.

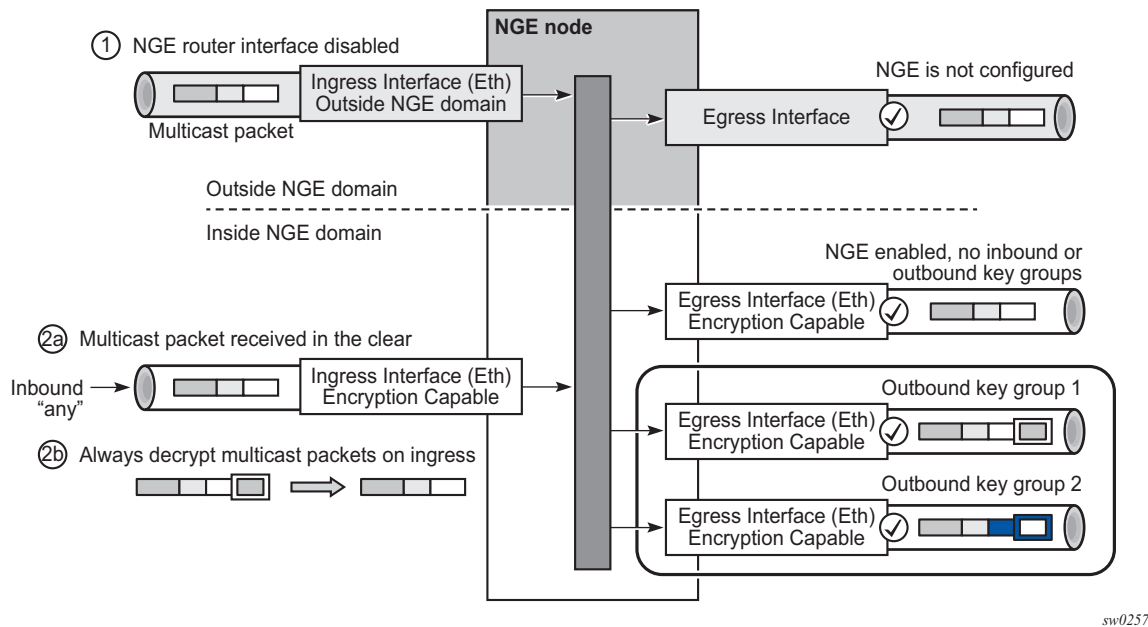
OSPFv3 IPsec support also uses IPsec transport mode packets. These packets originate from the CPM, which is considered outside the NGE domain; however, the above rules for encapsulating the packets with an NGE ESP apply and allow these packets to successfully transit the NGE domain.

2.7.7 Multicast Packets Traversing the NGE Domain

Multicast packets that traverse an NGE domain can be categorized into two main scenarios:

- Scenario 1 — multicast packets that ingress the router on an interface that is outside the NGE domain. These packets can egress a variety of interfaces that are either inside or outside the NGE domain.
- Scenario 2 — multicast packets that ingress the router on an interface that is inside the NGE domain. These packets can egress a variety of interfaces that are either inside or outside the NGE domain. This scenario has two cases:
 - Scenario 2a — the ingress multicast packet is not yet NGE-encrypted
 - Scenario 2b — the ingress multicast packet is NGE-encrypted

[Figure 22](#) shows these scenarios.

Figure 22 Processing Multicast Packets

Multicast packets received from outside the NGE domain (Scenario 1) are processed similarly to multicast packets received from inside the NGE domain (Scenarios 2a and 2b).

The processing rule is that multicast packets are always forwarded as clear text over the fabric. This means that for Scenario 2b, when a multicast packet is received on an encryption-capable interface and is NGE-encrypted, the packet is always decrypted first so that it can be processed in the same way as packets in Scenarios 1 and 2a.

On egress, the following scenarios apply:

- egressing an interface outside the NGE domain — packets are processed in the same way as any multicast packets forwarded out a non-NGE interface
- egressing an NGE router interface and no inbound or outbound key group is configured — the router forwards these packets out from the egress interface without encrypting them since there is no outbound key group configured. This behavior also applies to unicast packets in the same scenario.
- egressing an NGE router interface with the outbound key group configured — the router encrypts the multicast packet using the SPI keys of the outgoing SA configured in the key group. This behavior also applies to unicast packets in the same scenario.

2.7.8 Assigning Key Groups to Router Interfaces

Assigning key groups to router interfaces involves the following three steps:

Step 1. Enable NGE with the **group-encryption** command.

Step 2. Configure the outbound key group.

Step 3. Configure the inbound key group.

Step 1 is required so that the router can initialize and differentiate the interface for NGE traffic before accepting or sending NGE packets. This assigns the interface to an NGE domain.

Assigning key groups to a router interface in steps 2 and 3 is similar to assigning key groups to SDPs or VPRN-based services. An outbound key group cannot be configured for a router interface without first enabling **group-encryption**.

When group-encryption is enabled and no inbound key group is configured, the router will accept NGE Layer 3 packets that were encrypted using keys from any security association configured in any key group on the system. If the packet specifies a security association that is not configured in any key group on the node, the packet is dropped.

The outbound key group references the key group to use when traffic egresses the router on the router interface. The inbound key group is used to make sure ingress traffic is using the correct key group on the router interface. If ingress traffic is not using the correct key group, the router counts these packets as errors.

2.7.9 NGE and BFD Support

When NGE is enabled on a router interface, BFD packets that originate from the network processor on the adapter card or from the system are encrypted in the same way as BFD packets that are generated by the CPM.

2.7.10 NGE and ACL Interactions

When NGE is enabled on a router interface, the ACL function is applied as follows:

- on ingress — Normal ACLs are applied to traffic received on the interface that could be either NGE-encrypted or clear text. For NGE-encrypted packets, this implies that only the source, destination, and IP options are available to filter on ingress, as the protocol is ESP and the packet is encrypted. If an IP exception ACL is also configured on the interface, the IP exception ACL is applied first to allow any clear text packets to ingress as needed. After the IP exception ACL is applied and if another filter or ACL is configured on the interface, the other filter will process the remaining packet stream (NGE-encrypted and IP exception ACL packets), and other ACL functions such as PBR or Layer 4 information filtering could be applied to any clear text packets that passed the exception ACL.
- on egress — ACLs are applied to packets before they are NGE-encrypted as per normal operation without NGE enabled.

2.7.11 Router Interface NGE and ICMP Interactions Over the NGE Domain

Typically, ICMP works as expected over an NGE domain when all routers participating in the NGE domain are NGE-capable; this includes running an NGE domain over a private IP/MPLS network. When an ICMP message is required, the NGE packet is decrypted first and the original packet is restored to create a detailed ICMP message using the original packet's header information.

When the NGE domain crosses a Layer 3 service provider, or crosses over routers that are not NGE-aware, it is not possible to create a detailed ICMP message using the original packet's information, as the NGE packet protocol is always set to ESP. Furthermore, the NGE router that receives these ICMP messages will drop them because the messages are not NGE-encrypted.

The combination of dropping ICMP messages at the NGE border node and the missing unencrypted packet details in the ICMP information can cause problems with diagnosing network issues.

To help with diagnosing network issues, additional statistics are available on the interface to show whether ICMP messages are being returned from a foreign node. The following statistics are included in the group encryption NGE statistics for an interface:

- Group Enc Rx ICMP DestUnRch Pkts
- Group Enc Rx ICMP TimeExc Pkts
- Group Enc Rx ICMP Other Pkts

These statistics are used when clear text ICMP messages are received on an NGE router interface. The Invalid ESP statistics are not used in this situation even though the packet does not have a correct NGE ESP header. If there is no ingress exception ACL configured on the interface to allow the ICMP messages to be forwarded, the messages are counted and dropped.

If more information is required for these ICMP messages, such as source or destination address information, a second ICMP filter can be configured on the interface to allow logging of the ICMP messages. If the original packet information is also required, an egress exception ACL can be configured with the respective source or destination address information, or other criteria, to allow the original packet to enter the NGE domain in clear text and determine which flows are causing the ICMP failures.

2.7.12 1588v2 Encryption With NGE

If a router interface is enabled for encryption and Layer 3 1588v2 packets are sent, they will be encrypted using NGE. This means that if port timestamping is enabled on a router interface with NGE, the port timestamp is applied to the Layer 3 1588v2 packet using software-based timestamping instead of hardware-based timestamping, and consequently, timing accuracy may degrade. The exact level of timing or synchronization degradation is dependent on many factors, and testing is recommended to measure any impact.

If there is a need to support Layer 3 1588v2 with better accuracy for frequency or better time using port timestamping, an NGE exception ACL is required to keep the Layer 3 1588v2 packets in clear text. The exception ACL must enable UDP packets with destination port 319 to be sent in clear text.

2.8 Process Overview

The following items are components to configure basic router parameters:

- **Interface** — A logical IP routing interface. When created, attributes like an IP address, port, link aggregation group, or the system can be associated with the IP interface.
- **Address** — The address associates the device's system name with the IP system address. An IP address must be assigned to each IP interface.
- **System interface** — This creates an association between the logical IP interface and the system (loopback) address. The system interface address is the circuitless address (loopback) and is used by default as the router ID for protocols such as OSPF and BGP.
- **Router ID** — (Optional) The router ID specifies the router's IP address.
- **Autonomous system** — (Optional) An autonomous system (AS) is a collection of networks that are subdivided into smaller, more manageable areas.
- **Confederation** — (Optional) Creates confederation-autonomous systems within an AS to reduce the number of IBGP sessions required within an AS.

2.9 Configuration Notes

The following information describes router configuration requirements:

- A system interface and associated IP address must be specified.
- Boot options file (BOF) parameters must be configured before configuring router parameters.
- Confederations can be configured before protocol connections (such as BGP) and peering parameters are configured.
- IPv6 interfaces and associated routing protocols may only be configured on the following systems:
 - 7950 XRS systems
 - 7750 SR chassis systems
 - 7750 SR-a chassis systems
 - 7750 SR-e chassis systems
 - 7450 ESS systems running in mixed-mode with IPv6 functionality limited to those interfaces on slots with 7750 IOM3-XP/IMMs (or later) line cards.
 - 7750 SR-c4/12.

2.10 Configuring an IP Router with CLI

This section provides information to configure an IP router using CLI.

2.10.1 Router Configuration Overview

In a Nokia router, an interface is a logical named entity. An interface is created by specifying an interface name under the **config>router** context. This is the global router configuration context where objects like static routes are defined. An IP interface name can be up to 32 alphanumeric characters, must start with a letter, and is case-sensitive; for example, the interface name “1.1.1.1” is not allowed, but “int-1.1.1.1” is allowed.

To create an interface, the basic configuration tasks that must be performed are:

- Assign a name to the interface.
- Associate an IP address with the interface.
- Associate the interface with a network interface or the system interface.
- Configure appropriate routing protocols.

A system interface and network interface must be configured.

2.10.1.1 System Interface

The system interface is associated with a network entity (such as a specific Nokia router), not a specific interface. The system interface is also referred to as the loopback address. The system interface is associated during the configuration of the following entities:

- Termination point of service tunnels
- Hops when configuring MPLS paths and LSPs
- Addresses on a target router for BGP and LDP peering

The system interface is used to preserve connectivity (when routing reconvergence is possible) when an interface fails or is removed. The system interface is used as the router identifier. A system interface must have an IP address with a 32-bit subnet mask.

2.10.1.2 Network Interface

A network interface can be configured on one of the following entities:

- Physical or logical port
- SONET/SDH channel

For the 7950 XRS, a network interface can be configured on either a physical port or Ethernet LAG interface.

2.10.2 Basic Configuration

Refer to each specific chapter for specific routing protocol information and command syntax to configure protocols such as OSPF and BGP.

The most basic router configuration must have the following:

- System name
- System address

The following example shows a router configuration for the 7750 SR and 7450 ESS:

```
A:ALA-A> config# info
. . .
#-----
# Router Configuration
#-----
    router
        interface "system"
            address 10.10.10.103/32
        exit
        interface "to-104"
            address 10.0.0.103/24
            port 1/1/1
        exit
        exit
        autonomous-system 100
        confederation 1000 members 100 200 300
        router-id 10.10.10.103
    ...
        exit
        isis
        exit
    ...
#-----
A:ALA-A> config#
```

2.10.3 Common Configuration Tasks

The following sections describe basic system tasks.

2.10.3.1 Configuring a System Name

Use the system command to configure a name for the device. The name is used in the prompt string. Only one system name can be configured. If multiple system names are configured, the last one configured will overwrite the previous entry.

If special characters are included in the system name string, such as spaces, #, or ?, the entire string must be enclosed in double quotes. To configure the system name:

CLI Syntax: config# system
 name *system-name*

Example: config# system
 config>system# name ALA-A
 ALA-A>config>system# exit all
 ALA-A#

The following example shows the system name output:

```
A:ALA-A>config>system# info
#-----
# System Configuration
#-----
      name "ALA-A"
      location "Mt.View, CA, NE corner of FERG 1 Building"
      coordinates "37.390, -122.05500 degrees lat."
      snmp
      exit
```

2.10.3.2 Configuring Interfaces

The following command sequences create a system and a logical IP interface. The system interface assigns an IP address to the interface, then associates the IP interface with a physical port. The logical interface can associate attributes like an IP address or port.

The system interface cannot be deleted.

2.10.3.2.1 Configuring a System Interface

To configure a system interface:

CLI Syntax:

```
config>router
      interface interface-name
        address {ip-address/mask | ip-address
                  [netmask] } [broadcast {all-ones | host-ones}]
        secondary { [address/mask | ip-address]
                    [netmask] } [broadcast {all-ones | host-
                    ones}] [igp-inhibit]
```

2.10.3.2.2 Configuring a Network Interface

To configure a network interface for the 7450 ESS:

CLI Syntax:

```
config>router
      interface interface-name
        address ip-addr{/mask-length | mask}
          [broadcast {all-ones | host-ones}]
        cflowd {acl | interface}
        egress
          filter ip ip-filter-id
        ingress
          filter ip ip-filter-id
        port port-name
```

To configure a network interface for the 7750 SR:

CLI Syntax:

```
config>router
      interface interface-name
        address ip-addr{/mask-length | mask}
          [broadcast {all-ones | host-ones}]
        cflowd {acl | interface}
        egress
          filter ip ip-filter-id
          filter ipv6 ipv6-filter-id
        ingress
          filter ip ip-filter-id
          filter ipv6 ipv6-filter-id
        port port-name
```

To configure a network interface on the 7950 XRS:

CLI Syntax:

```
config>router
      interface interface-name
```

```

address ip-addr{/mask-length | mask}
    [broadcast {all-ones | host-ones}]
egress
    filter ip ip-filter-id
    filter ipv6 ipv6-filter-id
ingress
    filter ip ip-filter-id
    filter ipv6 ipv6-filter-id
port port-name

```

The following shows interface information about an IP configuration:

```

A:ALA-A>config>router# info
#-----
# IP Configuration
#-----
    interface "system"
        address 10.10.0.4/32
    exit
    interface "to-ALA-2"
        address 10.10.24.4/24
        port 1/1/1
        egress
            filter ip 10
        exit
    exit
...
#-----
A:ALA-A>config>router#

```

To enable CPU protection:

CLI Syntax: config>router
 interface interface-name
 cpu-protection policy-id

CPU protection policies are configured in the **config>sys>security>cpu-protection** context. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR System Management Guide*.

2.10.3.2.3 Assigning a Key Group to a Router Interface

Use the following CLI syntax to assign a key group to a router interface:

CLI Syntax: config>router# interface ip-int-name [create]
 group-encryption
 encryption-keygroup keygroup-id direction
 {inbound | outbound}

The following example displays a key group assigned to a router interface:

Example:

```
config>router# interface demo
config>router>if# group-encryption
config>router>if>group-encryp# encryption-keygroup 6
direction inbound
config>router>if>group-encryp# encryption-keygroup 6
direction outbound
```

The following example displays key group configuration for a router interface.

```
domain1>config>router# info
-----
...
    interface demo
        group-encryption
            encryption-keygroup 6 direction inbound
            encryption-keygroup 6 direction outbound
        exit
        no shutdown
        exit
    exit
...
-----
```

2.10.3.2.4 Configuring IPv6 Parameters

IPv6 interfaces and associated routing protocols may only be configured on the following systems:

- 7950 XRS systems.
- 7750 SR chassis systems.
- 7750 SR-a chassis systems.
- 7750 SR-e chassis systems.
- 7450 ESS chassis running in mixed-mode, with IPv6 functionality limited to those interfaces on slots with 7750 IOM3-XP/IMMs (or later) line card.
- 7750 SR-c4/12.

The following displays the interface configuration showing the IPv6 default configuration when IPv6 is enabled on the interface:

```
A:ALA-49>config>router>if>ipv6# info detail
-----
` port 1/2/37
  ipv6
    packet-too-big 100 10
    param-problem 100 10
    redirects 100 10
```

```

        time-exceeded 100 10
        unreachable 100 10
    exit
-----
A:ALA-49>config>router>if>ipv6# exit all

```

To configure IPv6 parameters on a router interface:

CLI Syntax:

```

config>router# interface interface-name
port port-name
ipv6
    address {ipv6-address/prefix-length} [eui-64]
    icmp6
        packet-too-big [number seconds]
        param-problem [number seconds]
        redirects [number seconds]
        time-exceeded [number seconds]
        unreachable [number seconds]
    neighbor ipv6-address mac-address

```

The following displays a configuration example showing interface information:

```

A:ALA-49>config>router>if# info
-----
    address 10.11.10.1/24
    port 1/2/37
    ipv6
        address 10::1/24
    exit
-----
A:ALA-49>config>router>if#

```

2.10.3.2.5 Configuring IPv6 Over IPv4 Parameters

The following sections provide several examples of the features that must be configured (tunnel ingress and egress node) to implement IPv6 over IPv4 relay services for the 7750 SR OS.

2.10.3.2.6 Tunnel Ingress Node

The following example shows the configuration of the interface through which the IPv6 over IPv4 traffic leaves the node. This must be configured on a network interface.

CLI Syntax:

```

config>router
    static-route-entry::C8C8:C802/128
        indirect 200.200.200.2

```

```
interface ip-int-name
    address {ip-address/mask | ip-address
            netmask} [broadcast {all-ones |
            host-ones}]
    port port-name
```

The following example shows an interface configuration:

```
A:ALA-49>config>router# info
-----
...
    interface "ip-1.1.1.1"
        address 1.1.1.1/30
        port 1/1/1
    exit
...
-----
A:ALA-49>config>router#
```

Both the IPv4 and IPv6 system addresses must be configured:

CLI Syntax:

```
config>router
    interface ip-int-name
        address {ip-address/mask | ip-address netmask}
        [broadcast {all-ones | host-ones}]
        ipv6
            address ipv6-address/prefix-length [eui-
            64]
```

The following example shows the configuration of interface information:

```
A:ALA-49>config>router# info
-----
...
    interface "system"
        address 200.200.200.1/32
        ipv6
            address 3FFE::C8C8:C801/128
        exit
    exit
...
-----
A:ALA-49>config>router#
```

Learning the Tunnel Endpoint IPv4 System Address

The following example shows the OSPF configuration to learn the IPv4 system address of the tunnel endpoint:

CLI Syntax:

```
config>router
ospf
    area area-id
        interface ip-int-name
```

The following example shows the configuration of OSPF output:

```
A:ALA-49>config>router# info
-----
...
    ospf
        area 0.0.0.0
            interface "system"
            exit
            interface "ip-1.1.1.1"
            exit
        exit
    exit
-----
A:ALA-49>config>router#
```

Configuring IPv4 BGP Peer

The following example shows the configuration of an IPv4 BGP peer with (IPv4 and) IPv6 protocol families:

CLI Syntax:

```
config>router
bgp
    export policy-name [policy-name...(upto 5 max)]
    router-id ip-address
    group name
        family [ipv4] [vpn-ipv4] [ipv6] [mcast-ipv4]
        type {internal | external}
        neighbor ip-address
            local-as as-number [private]
            peer-as as-number
```

The following example shows the configuration of BGP output:

```
A:ALA-49>config>router# info
-----
...
    bgp
        export "ospf3"
        router-id 200.200.200.1
        group "main"
            family ipv4 ipv6
            type internal
            neighbor 200.200.200.2
                local-as 1
                peer-as 1
```

```

        exit
    exit
exit
...
-----
A:ALA-49>config>router#

```

An Example of a IPv6 Over IPv4 Tunnel Configuration

The IPv6 address is the next-hop as it is received through BGP. The IPv4 address is the system address of the tunnel's endpoint.

The following example shows the configuration of a policy to export IPv6 routes into BGP:

CLI Syntax:

```

config>router
bgp
    export policy-name [policy-name...(upto 5 max)]
    router-id ip-address
    group name
        family [ipv4] [vpn-ipv4] [ipv6] [mcast-ipv4]
        type {internal | external}
        neighbor ip-address
            local-as as-number [private]
            peer-as as-number

```

The following example shows the configuration output:

```

A:ALA-49>config>router# info
-----
...
    policy-options
        policy-statement "ospf3"
            description "Plcy Stmnt For 'From ospf3 To bgp'"
            entry 10
                description "Entry From Protocol ospf3 To bgp"
                from
                    protocol ospf3
                exit
                to
                    protocol bgp
                exit
                action accept
            exit
        exit
    exit
exit
...
-----
A:ALA-49>config>router#

```

2.10.3.2.7 Tunnel Egress Node

The following example shows the configuration of the interface through which the IPv6 over IPv4 traffic leaves the node. It must be configured on a network interface. Both the IPv4 and IPv6 system addresses must be configured.

CLI Syntax:

```
config>router
    static-route ::C8C8:C801/128
        indirect 200.200.200.1
            interface ip-int-name
                address {ip-address/mask> | ip-
                    address netmask} [broadcast
                        {all-ones | host-ones}]
                ipv6
                address ipv6-address/prefix-length
                    [eui-64]
                port port-name
```

The following example shows the interface configuration:

```
A:ALA-49>config>router# info
-----
...
    interface "ip-1.1.1.2"
        address 1.1.1.2/30
        port 1/1/1
    exit
    interface "system"
        address 200.200.200.2/32
        ipv6
            address 3FFE::C8C8:C802/128
        exit
    exit
-----
```

Learning the Tunnel Endpoint IPv4 System Address

The following example shows the configuration of the OSPF configuration to learn the IPv4 system address of the tunnel endpoint:

CLI Syntax:

```
config>router
    ospf
        area area-id
            interface ip-int-name
```

The following example shows the configuration of OSPF information:

```
A:ALA-49>config>router# info
-----
```

```

...
    ospf
      area 0.0.0.0
        interface "system"
        exit
        interface "ip-1.1.1.2"
        exit
      exit
    exit
  -----
A:ALA-49>config>router#

```

Configuring an IPv4 BGP Peer

The following example shows the configuration an IPv4 BGP peer with (IPv4 and) IPv6 protocol families:

CLI Syntax:

```

config>router
bgp
  export policy-name [policy-name...(upto 5 max)]
  router-id ip-address
  group name
    family [ipv4] [vpn-ipv4] [ipv6] [mcast-ipv4]
    type {internal | external}
    neighbor ip-address
      local-as as-number [private]
      peer-as as-number

```

The following example shows the IPv4 BGP peer configuration:

```

A:ALA-49>config>router# info
-----
...
  bgp
    export "ospf3"
    router-id 200.200.200.2
    group "main"
      family ipv4 ipv6
      type internal
      neighbor 200.200.200.1
        local-as 1
        peer-as 1
      exit
    exit
  exit
...
-----
A:ALA-49>config>router#

```

An Example of a IPv6 Over IPv4 Tunnel Configuration

The IPv6 address is the next-hop as it is received through BGP. The IPv4 address is the system address of the tunnel's endpoint.

The following example shows the configuration of a policy to export IPv6 routes into BGP:

CLI Syntax:

```
config>router
bgp
    export policy-name [policy-name...(upto 5 max)]
    router-id ip-address
    group name
        family [ipv4] [vpn-ipv4] [ipv6] [mcast-ipv4]
        type {internal | external}
        neighbor ip-address
            local-as as-number [private]
            peer-as as-number
```

The following example shows an IPv6 over IPv4 tunnel configuration:

```
A:ALA-49>config>router# info
-----
...
    policy-options
        policy-statement "ospf3"
            description "Plcy Stmtnt For 'From ospf3 To bgp'"
            entry 10
                description "Entry From Protocol ospf3 To bgp"
                from
                    protocol ospf3
                exit
                to
                    protocol bgp
                exit
                action accept
                exit
            exit
        exit
    exit
exit
-----
A:ALA-49>config>router#
```

2.10.3.2.8 Router Advertisement

To configure the router to originate router advertisement messages on an interface, the interface must be configured under the **router-advertisement** context and be enabled (no shutdown). All other router advertisement configuration parameters are optional.

Router advertisement can be configured under the **config>router>router-advertisement** context or under the **config>service>vprn>router-advertisement** context. Use the following examples of CLI syntax to enable router advertisement and configure router advertisement parameters.

To configure router advertisement on the 7750 SR:

CLI Syntax:

```
config>router# router-advertisement
dns-options
    rdnss-lifetime seconds
    dns-servers ipv6-address
interface ip-int-name
    current-hop-limit number
    dns-options
        rdnss-lifetime {seconds | infinite}
        dns-servers ipv6-address
    include-dns
    managed-configuration
    max-advertisement-interval seconds
    min-advertisement-interval seconds
    mtu mtu-bytes
    other-stateful-configuration
    prefix ipv6-prefix/prefix-length
        autonomous
        on-link
        preferred-lifetime {seconds | infinite}
        valid-lifetime {seconds | infinite}
    reachable-time milliseconds
    retransmit-time milliseconds
    router-lifetime seconds
    no shutdown
    use-virtual-mac
```

To configure router advertisement for the 7450 ESS:

CLI Syntax:

```
config>router# router-advertisement
dns-options
    rdnss-lifetime seconds
interface ip-int-name
    current-hop-limit number
    dns-options
        rdnss-lifetime {seconds | infinite}
    include-dns
    managed-configuration
    max-advertisement-interval seconds
    min-advertisement-interval seconds
    mtu mtu-bytes
    other-stateful-configuration
        autonomous
```

```

on-link
preferred-lifetime {seconds | infinite}
valid-lifetime {seconds | infinite}
reachable-time milliseconds
retransmit-time milliseconds
router-lifetime seconds
no shutdown
use-virtual-mac

```

The following example shows a router advertisement configuration:

```

*A:sim131>config>router>router-advert# info
-----
        interface "n1"
            prefix 2001:db8:3::/64
            exit
            use-virtual-mac
            no shutdown
        exit
-----
*A:sim131>config>router>router-advert# interface n1
*A:sim131>config>router>router-advert>if# prefix 2001:db8:3::/64
-----
            autonomous
            on-link
            preferred-lifetime 604800
            valid-lifetime 2592000
-----
*A:tahi>config>router>router-advert>if>prefix#

```

2.10.3.2.9 Configuring IPv6 Parameters

The following example shows the IPv6 default configuration when IPv6 is enabled on the interface:

```

A:ALA-49>config>router>if>ipv6# info detail
-----
    port 1/3/37
    ipv6
        packet-too-big 100 10
        param-problem 100 10
        redirects 100 10
        time-exceeded 100 10
        unreachablees 100 10
    exit
-----
A:ALA-49>config>router>if>ipv6# exit all

```

The following example shows an IPv6 configuration:

```

A:ALA-49>config>router>if# info

```

```

-----
        address 10.11.10.1/24
        port 1/3/37
        ipv6
            address 10::1/24
        exit
-----
A:ALA-49>config>router>if#

```

An Example of a IPv6 Over IPv4 Tunnel Configuration

The IPv6 address is the next-hop as it is received through BGP. The IPv4 address is the system address of the tunnel's endpoint.

The following example shows the configuration of a policy to export IPv6 routes into BGP:

CLI Syntax:

```

config>router
bgp
    export policy-name [policy-name...(upto 5 max)]
    router-id ip-address
    group name
        family [ipv4] [vpn-ipv4] [ipv6] [mcast-ipv4]
        type {internal | external}
        neighbor ip-address
            local-as as-number [private]
            peer-as as-number

```

The following example shows the configuration of the policy output:

```

A:ALA-49>config>router# info
-----
...
    policy-options
        policy-statement "ospf3"
            description "Plcy Stmt For 'From ospf3 To bgp'"
            entry 10
                description "Entry From Protocol ospf3 To bgp"
                from
                    protocol ospf3
                exit
                to
                    protocol bgp
                exit
                action accept
                exit
            exit
        exit
    exit
-----
A:ALA-49>config>router#

```


2.10.3.2.10 Configuring Proxy ARP

To configure proxy ARP, you can configure:

- A prefix list in the **config>router>policy-options>prefix-list** context.
- A route policy statement in the **config>router>policy-options>policy-statement** context and apply the specified prefix list.
 - In the policy statement **entry>to** context, specify the host source address(es) for which ARP requests can or cannot be forwarded to non-local networks, depending on the specified action.
 - In the policy statement **entry>from** context, specify network prefixes that ARP requests will or will not be forwarded to depending on the action if a match is found. For more information about route policies, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Unicast Routing Protocols Guide*.
- Apply the policy statement to the **proxy-arp** configuration in the **config>router>interface** context.

CLI Syntax:

```
config>router# policy-options
begin
commit
prefix-list name
        prefix ip-prefix/mask [exact | longer | through
        length | prefix-length-range length1-length2]
```

To configure the policy statement specified in the **proxy-arp-policy policy-statement** command:

CLI Syntax:

```
config>router# policy-options
begin
commit
policy-statement name
        default-action {accept | next-entry | next-policy |
        reject}
        entry entry-id
                action {accept | next-entry | next-policy |
                reject}
        to
                prefix-list name [name...(upto 5 max)]
        from
                prefix-list name [name...(upto 5 max)]
```

The following example shows the prefix list and policy statement configuration:

```
A:ALA-49>config>router>policy-options# info
-----
        prefix-list "prefixlist1"
```

```

        prefix 10.20.30.0/24 through 32
    exit
    prefix-list "prefixlist2"
        prefix 10.10.10.0/24 through 32
    exit
...
    policy-statement "ProxyARPolicy"
        entry 10
            from
                prefix-list "prefixlist1"
            exit
            to
                prefix-list "prefixlist2"
            exit
            action reject
        exit
        default-action accept
    exit
exit
...
-----
A:ALA-49>config>router>policy-options#

```

Use the following CLI to configure proxy ARP:

CLI Syntax: config>router>interface *interface-name*
 local-proxy-arp
 proxy-arp-policy *policy-name* [*policy-name...*(upto 5
 max)]
 remote-proxy-arp

The following example shows a proxy ARP configuration:

```

A:ALA-49>config>router>if# info
-----
        address 128.251.10.59/24
        local-proxy-arp
        proxy-arp
            policy-statement "ProxyARPolicy"
        exit
-----
A:ALA-49>config>router>if#

```

2.10.3.2.11 Creating an IP Address Range

An IP address range can be reserved for exclusive use for services by defining the **config>router>service-prefix** command. When the service is configured, the IP address must be in the range specified as a service prefix. If no service prefix command is configured, no limitation exists.

The **no service-prefix** *ip-prefix/mask* command removes all address reservations. A service prefix cannot be removed while one or more services use address(es) in the range to be removed.

CLI Syntax: `config>router`
 `service-prefix ip-prefix/mask [exclusive]`

2.10.3.3 Deriving the Router ID

The router ID defaults to the address specified in the system interface command. If the system interface is not configured with an IP address, the router ID inherits the last four bytes of the MAC address. The router ID can also be manually configured in the **config>router** *router-id* context. On the BGP protocol level, a BGP router ID can be defined in the **config>router>bgp** *router-id* context and is only used within BGP.

If a new router ID is configured, protocols are not automatically restarted with the new router ID. The next time a protocol is initialized, the new router ID is used. An interim period of time can occur when different protocols use different router IDs. To force the new router ID, issue the shutdown and no shutdown commands for each protocol that uses the router ID, or restart the entire router.

It is possible to configure SR OS to operate with an IPv6 only BOF and no IPv4 system interface address. When configured in this manner, the operator must explicitly define IPv4 router IDs for protocols such as OSPF and BGP because there is no mechanism to derive the router ID from an IPv6 system interface address.

To configure the router ID:

CLI Syntax: `config>router`
 `router-id router-id`
 `interface ip-int-name`
 `address {ip-address/mask | ip-address netmask}`
 `[broadcast {all-ones | host-ones}]`

The following example shows a router ID configuration:

```
A:ALA-4>config>router# info
#-----
# IP Configuration
#-----
      interface "system"
        address 10.10.0.4/32
      exit
      . . .
      router-id 10.10.0.4
#-----
```

```
A:ALA-4>config>router#
```

2.10.3.4 Configuring a Confederation

Configuring a confederation is optional. The AS and confederation topology design should be carefully planned. Autonomous system (AS), confederation, and BGP connection and peering parameters must be explicitly created on each participating router. Identify AS numbers, confederation numbers, and members participating in the confederation.

See the BGP section for CLI syntax and command descriptions.

To configure a confederation:

CLI Syntax: `config>router`
 `confederation confed-as-num members member-as-num`

The following example shows the configuration of the confederation topology in [Confederation Configuration](#).



Note:

- Confederations can be preconfigured prior to configuring BGP connections and peering.
- Each confederation can have up to 15 members.

The following example shows a confederation configuration:

```
A:ALA-B>config>router# info
#-----
# IP Configuration
#-----
      interface "system"
        address 10.10.10.103/32
      exit
      interface "to-104"
        shutdown
        address 10.0.0.103/24
        port 1/1/1
      exit
      autonomous-system 100
      confederation 2002 members 200 300 400
      router-id 10.10.10.103

#-----
A:ALA-B>config>router#
```

2.10.3.5 Configuring an Autonomous System

Configuring an autonomous system is optional. To configure an autonomous system:

CLI Syntax: `config>router`
 `autonomous-system as-number`

The following example shows an autonomous system configuration:

```
A:ALA-A>config>router# info
#-----
# IP Configuration
#-----
      interface "system"
        address 10.10.10.103/32
      exit
    interface "to-104"
      address 10.0.0.103/24
      port 1/1/1
      exit
    exit
    autonomous-system 100
    router-id 10.10.10.103
#-----
A:ALA-A>config>router#
```

2.10.3.6 Configuring Overload State on a Single SFM

When a router has fewer than the full set of SFMs functioning, the forwarding capacity can be reduced. Some scenarios include:

- fewer than the maximum number of SFMs installed in the system
- one or more SFMs have failed
- the system is in the ISSU process and the SFM is co-located on the CPM

An overload condition can be set for IS-IS and OSPF to enable the router to still participate in exchanging routing information, but route all traffic away from it when insufficient SFMs are active. This is achieved using the following CLI commands:

CLI Syntax: `config>router>single-sfm-overload [holdoff-time hold-off-time]`
 `config>service>vprn>single-sfm-overload [holdoff-time hold-off-time]`
 `tools>perform>redundancy>forced-single-sfm-overload`

These cause an overload state in the IGP to trigger the traffic reroute by setting the overload bit in IS-IS or setting the metric to maximum in OSPF. When PIM uses IS-IS or OSPF to find out the upstream router, a next-hop change in the IS-IS or OSPF will cause PIM to join the new path and prune the old path, which effectively also reroutes the multicast traffic downstream as well as the unicast traffic.

When the problem is resolved, and the required compliment of SFMs become active in the router, the overload condition is cleared, which will cause the traffic to be routed back to the router.

The conditions to set overload are:

- 7750 SR-12/SR-7/SR-c12 and 7450 ESS-12/ESS-7/ESS-6 platforms: protocol sets overload if one of the SF/CPMs fails
- 7750 SR-12e and 7950 XRS platforms: protocol sets overload if two SFMs fail (two SFMs belonging to different SFM pairs on the XRS-40)

2.11 Service Management Tasks

This section describes IP router service management tasks:

2.11.1 Changing the System Name

The system command sets the name of the device and is used in the prompt string. Only one system name can be configured. If multiple system names are configured, the last one configured will overwrite the previous entry.

To change the system name:

CLI Syntax: `config# system`
 `name system-name`

The following example shows the configuration to change the system name:

Example: `A:ALA-A>config>system# name tgif`
 `A:TGIF>config>system#`

The following example shows the system name change:

```
A:ALA-A>config>system# name TGIF
A:TGIF>config>system# info
#-----
# System Configuration
#-----
      name "TGIF"
      location "Mt.View, CA, NE corner of FERG 1 Building"
      coordinates "37.390, -122.05500 degrees lat."
      synchronize
      snmp
          exit
          security
              snmp
                  community "private" rwa version both
              exit
          exit
          . . .
#-----
A:TGIF>config>system#
```

2.11.2 Modifying Interface Parameters

Starting at the **config>router** level, navigate down to the **router interface** context.

To modify an IP address:

CLI Syntax:

```
A:ALA-A>config>router# interface "to-sr1"
A:ALA-A>config>router>if# shutdown
A:ALA-A>config>router>if# no address
A:ALA-A>config>router>if# address 10.0.0.25/24
A:ALA-A>config>router>if# no shutdown
```

To modify a port:

CLI Syntax:

```
A:ALA-A>config>router# interface "to-sr1"
A:ALA-A>config>router>if# shutdown
A:ALA-A>config>router>if# no port
A:ALA-A>config>router>if# port 1/1/2
A:ALA-A>config>router>if# no shutdown
```

The following example shows the interface configuration:

```
A:ALA-A>config>router# info
#-----
# IP Configuration
#-----
      interface "system"
        address 10.0.0.103/32
      exit
      interface "to-sr1"
        address 10.0.0.25/24
        port 1/1/2
      exit
      router-id 10.10.0.3
#-----
A:ALA-A>config>router#
```

2.11.3 Removing a Key Group from a Router Interface

Use the following CLI syntax to remove a key group from a router interface:

CLI Syntax:

```
config>router# interface ip-int-name
                  group-encryption
                    no encryption-keygroup keygroup-id direction
                      {inbound | outbound}
```

The following example displays a key group removed from a router interface:

Example:

```
config>router# interface demo
config>router>if# group-encryption
config>router>if>group-encryp# no encryption-keygroup 6
                        direction inbound
```



```
config>router>if>group-encryp# no encryption-keygroup 6
direction outbound
```

The following example shows that the key group configuration has been removed from a router interface.

```
domain1>config>router# info
-----
...
    interface demo
        group-encryption
        exit
        no shutdown
        exit
    exit
...
-----
```

2.11.4 Changing the Key Group for a Router Interface

The following CLI syntax changes the key group on a router interface. In the example below, the inbound and outbound key groups are changed from key group 6 to key group 8.

CLI Syntax:

```
config>router# interface ip-int-name
group-encryption
    no encryption-keygroup keygroup-id direction
    {inbound | outbound}
```

Example:

```
config>router# interface demo
config>router>if# group-encryption
config>router>if>group-encryp# no encryption-keygroup 6
direction inbound
config>router>if>group-encryp# encryption-keygroup 8
direction outbound
config>router>if>group-encryp# encryption-keygroup 8
direction inbound
```

The following example shows that the key group configuration has been changed for the router interface.

```
domain1>config>router# info
-----
...
    interface demo
        group-encryption
            encryption-keygroup 8 direction inbound
            encryption-keygroup 8 direction outbound
    exit
...
-----
```

```
        exit
    no shutdown
    exit
exit
...
-----
```

2.11.5 Deleting a Logical IP Interface

The **no** form of the **interface** command typically removes the entry, but all entity associations must be shut down and/or deleted before an interface can be deleted.

- Step 1.** Before an IP interface can be deleted, it must first be administratively disabled with the shutdown command.
- Step 2.** After the interface has been shut down, it can then be deleted with the **no interface** command.

CLI Syntax:

```
config>router
no interface ip-int-name
```

Example:

```
config>router# interface test-interface
config>router>if# shutdown
config>router>if# exit
config>router# no interface test-interface
config>router#
```

2.12 IP Router Configuration Command Reference

- [Command Hierarchies](#)
- [Command Descriptions](#)

2.12.1 Command Hierarchies

- [Router Commands](#)
- [Router BFD Commands](#)
- [Router L2TP Commands](#)
- [Router Interface Commands](#)
- [Router Interface IPv6 Commands](#)
- [Router Advertisement Commands](#)

2.12.1.1 Router Commands

```

config
— [no] router [router-instance]
— aggregate ip-prefix/ip-prefix-length [summary-only] [as-set] [aggregator as-
  number:ip-address] [black-hole] [community comm-id] [description description]
— aggregate ip-prefix/ip-prefix-length [summary-only] [as-set] [aggregator as-
  number:ip-address] [community comm-id] [indirect ip-address] [description
  description]
— no aggregate ip-prefix/ip-prefix-length
— autonomous-system autonomous-system
— no autonomous-system
— [no] class-forwarding
— confederation confed-as-num members as-number [as-number...(up to 15 max)]
— no confederation [confed-as-num members as-number...(up to 15 max)]
— ecmp max-ecmp-routes
— no ecmp
— no entropy-label
— fib-priority {high | standard}
— flowspec
  — ip-filter-max-size {value | default}
  — ipv6-filter-max-size {value | default}
— [no] icmp-tunneling
— [no] ip-fast-reroute
— [no] ldp-shortcut
— mc-maximum-routes number [log-only] [threshold threshold]
— no mc-maximum-routes
— mpls-labels

```

- **bgp-labels-hold-timer** *seconds*
- **no bgp-labels-hold-timer**
- **sr-labels** **start** *start-value* **end** *end-value*
- **no sr-labels**
- **static-label-range** *static-range*
- **no static-label-range**
- **mss-adjust-group** *nat-group-id* **segment-size** *segment-size*
- **no mss-adjust-group**
- **multicast-info-policy** *policy-name*
- **no multicast-info-policy**
- **multicast-info-policy**
 - **description** *description-string*
 - **no description**
- **origin-validation**
 - **[no] rpki-session** *ip-address*
 - **[no] connect-retry** *seconds*
 - **[no] description** *description-string*
 - **[no] local-address** *ip-address*
 - **[no] port** *port-id*
 - **[no] refresh-time** *seconds* **hold-time** *seconds*
 - **[no] shutdown**
 - **[no] stale-time** *seconds*
 - **static-entry** *ip-prefix/prefix-length upto prefix-length2 origin-as as-number* [**valid** | **invalid**]
 - **no static-entry** *ip-prefix/prefix-length1-prefix-length2*
- **router-id** *ip-address*
- **no router-id**
- **service-prefix** {*ip-prefix/mask* | *ip-prefix netmask*} [**exclusive**]
- **no service-prefix** {*ip-prefix/mask* | *ip-prefix netmask*}
- **sgt-qos**
 - **application** *dscp-app-name* **dscp** {*dscp-value* | *dscp-name*}
 - **application** *dot1p-app-name* **dot1p** *dot1p-priority*
 - **no application** {*dscp-app-name* | *dot1p-app-name*}
 - **dscp** *dscp-name* **fc** *fc-name*
 - **[no] dscp** *dscp-name*
- **single-sfm-overload** [**holdoff-time** *holdoff-time*]
- **no single-sfm-overload**
- **[no] static-route-entry** {*ip-prefix/prefix-length*} [**mcast**]
 - **[no] black-hole**
 - **[no] community** *comm-id*
 - **[no] description** *description-string*
 - **[no] dynamic-bgp**
 - **[no] generate-icmp**
 - **[no] metric** *metric-value*
 - **[no] preference** *preference-value*
 - **[no] prefix-list** *name* {**all** | **none** | **any**}
 - **[no] shutdown**
 - **[no] tag** *tag-value*
 - **[no] indirect** *ip-address*
 - **[no] community** *comm-id*
 - **[no] cpe-check** *cpe-ip-address*
 - **[no] drop-count** *count*
 - **[no] interval** *seconds*
 - **[no] log**

```

    — [no] padding-size padding-size
  — [no] description description-string
  — [no] destination-class dest-index
  — [no] forwarding-class {be | l2 | af | l1 | h2 | ef | h1 | nc}
    — [no] priority {low | high}
  — [no] metric metric-value
  — [no] preference preference-value
  — [no] prefix-list prefix-list-name {all | none | any}
  — [no] shutdown
  — [no] source-class source-index
  — [no] tag tag-value
  — [no] tunnel-next-hop
    — [no] disallow-igp
    — [no] resolution
    — [no] resolution-filter {any | disable | filter}
      — [no] ldp
      — [no] rsvp-te
        — [no] lsp lsp-name
      — [no] sr-isis
      — [no] sr-ospf
      — [no] sr-te
        — [no] lsp lsp-name
  — [no] next-hop {ip-address | ip-int-name | ipv6 address}
    — [no] bfd-enable
    — [no] community comm-id
    — [no] cpe-check cpe-ip-address
      — [no] drop-count count
      — [no] interval seconds
      — [no] log
      — [no] padding-size padding-size
    — [no] description description-string
    — [no] destination-class dest-index
    — [no] forwarding-class {be | l2 | af | l1 | h2 | ef | h1 | nc}
      — [no] priority {low | high}
    — [no] ldp-sync
    — [no] metric metric-value
    — [no] preference preference-value
    — [no] prefix-list name {all | none | any}
    — [no] shutdown
    — [no] source-class [source-index]
    — [no] tag tag-value
    — [no] validate-next-hop
  — [no] triggered-policy
  — ttl-propagate
    — label-route-local [none | all]
    — label-route-transit [none | all]
    — lsr-label-route [none | all]
    — vprn-local [none | vc-only | all]
    — vprn-transit [none | vc-only | all]
  — weighted-ecmp

config
  — router management
    — origin-validation

```

- [no] **rpki-session** *ip-address*
 - [no] **connect-retry** *seconds*
 - [no] **description** *description-string*
 - [no] **local-address** *ip-address*
 - [no] **port** *port-id*
 - [no] **refresh-time** *seconds* **hold-time** *seconds*
 - [no] **shutdown**
 - [no] **stale-time** *seconds*

2.12.1.2 Router BFD Commands

- ```
config
— router
 — bfd
 — bfd-template name [create]
 — bfd-template name
 — transmit-interval transmit-interval
 — no transmit-interval
 — receive-interval receive-interval
 — no receive-interval
 — echo-receive echo-interval
 — no echo-receive
 — multiplier multiplier
 — no multiplier
 — [no] type cpm-np
```

### 2.12.1.3 Router L2TP Commands

The router L2TP commands apply only to the 7750 SR and 7450 ESS.

- ```
config
— router [router-name]
  — l2tp
    — calling-number-format ascii-spec
    — no calling-number-format
    — challenge {always}
    — no challenge
    — df-bit-lac {always | never}
    — no df-bit-lac
    — destruct-timeout destruct-timeout
    — no destruct-timeout
    — eth-tunnel
      — reconnect-timeout seconds
      — no reconnect-timeout
    — exclude-avps calling-number
    — no exclude-avps
    — group tunnel-group-name [create]
    — no group tunnel-group-name
```

-
- **avp-hiding** {sensitive | always}
 - **no avp-hiding**
 - **challenge** [always]
 - **no challenge**
 - **description** *description-string*
 - **no description**
 - **df-bit-lac** {always | never | default}
 - **no df-bit-lac**
 - **destruct-timeout** *destruct-timeout*
 - **no destruct-timeout**
 - **hello-interval** *hello-interval*
 - **no hello-interval**
 - **idle-timeout** *idle-timeout*
 - **no idle-timeout**
 - **l2tpv3**
 - **cookie-length** {4 | 8 | default}
 - **no cookie-length**
 - **digest-type** {default | md5 | sha1 | none}
 - **no digest-type**
 - **nonce-length** {*length* | default}
 - **no nonce-length**
 - **password** *password* [hash | hash2]
 - **no password**
 - **private-tcp-mss-adjust** *octets*
 - **private-tcp-mss-adjust** default
 - **no private-tcp-mss-adjust**
 - **public-tcp-mss-adjust** *octets*
 - **public-tcp-mss-adjust** default
 - **no public-tcp-mss-adjust**
 - **pw-cap-list** {ethernet | ethernet-vlan}
 - **no pw-cap-list**
 - **rem-router-id** *ip-addr*
 - **no rem-router-id**
 - [no] **track-password-change**
 - **lns-group** *lns-group-id*
 - **no lns-group**
 - **load-balance-method** {per-session | per-tunnel}
 - **no load-balance-method**
 - **local-address** *ip-address*
 - **no local-address**
 - **local-name** *host-name*
 - **no local-name**
 - **max-retries-estab** *max-retries*
 - **no max-retries-estab**
 - **max-retries-not-estab** *max-retries*
 - **no max-retries-not-estab**
 - **password** *password* [hash | hash2]
 - **no password**
 - **ppp**
 - **authentication** {chap | pap | pref-chap | pref-pap}
 - **authentication-policy** *auth-policy-name*
 - **no authentication-policy**
 - **default-group-interface** *ip-int-name* **service-id** *service-id*
 - **no default-group-interface**

-
- **keepalive** *seconds* [**hold-up-multiplier** *multiplier*]
 - **no keepalive**
 - **[no] lcp-force-ack-accm**
 - **mtu** *mtu-bytes*
 - **no mtu**
 - **[no] proxy-authentication**
 - **[no] proxy-lcp**
 - **user-db** *local-user-db-name*
 - **no user-db**
 - **session-assign-method** {**existing-first** | **weighted** | **weighted-random**}
 - **no session-assign-method**
 - **session-limit** **unlimited**
 - **session-limit** *session-limit*
 - **no session-limit**
 - **tunnel** *tunnel-name* [**create**]
 - **no tunnel** *tunnel-name*
 - **[no] auto-establish**
 - **avp-hiding** {**never** | **sensitive** | **always**}
 - **no avp-hiding**
 - **challenge** *challenge-mode*
 - **no challenge**
 - **description** *description-string*
 - **no description**
 - **df-bit-lac** {**always** | **never** | **default**}
 - **no df-bit-lac**
 - **destruct-timeout** *destruct-timeout*
 - **no destruct-timeout**
 - **hello-interval** *hello-interval*
 - **hello-interval** **infinite**
 - **no hello-interval**
 - **idle-timeout** *idle-timeout*
 - **idle-timeout** **infinite**
 - **no idle-timeout**
 - **load-balance-method** {**per-session** | **per-tunnel**}
 - **no load-balance-method**
 - **local-address** *ip-address*
 - **no local-address**
 - **local-name** *host-name*
 - **no local-name**
 - **max-retries-estab** *max-retries*
 - **no max-retries-estab**
 - **max-retries-not-estab** *max-retries*
 - **no max-retries-not-estab**
 - **password** *password* [**hash** | **hash2**]
 - **no password**
 - **peer** *ip-address*
 - **no peer**
 - **ppp**
 - **[no] lcp-force-ack-accm**
 - **preference** *preference*
 - **no preference**
 - **remote-name** *host-name*
 - **no remote-name**

```

— session-limit unlimited
— session-limit session-limit
— no session-limit
— [no] shutdown
— group-session-limit session-limit
— group-session-limit unlimited
— no group-session-limit
— l2tpv3
— cookie-length {4 | 8}
— no cookie-length
— digest-type {default | md5 | sha1 | none}
— no digest-type
— nonce-length length
— nonce-length default
— no nonce-length
— password password [hash | hash2]
— no password
— private-tcp-mss-adjust octets
— no private-tcp-mss-adjust
— public-tcp-mss-adjust octets
— no public-tcp-mss-adjust
— transport-type ip
— no transport-type
— next-attempt {same-preference-level | next-preference-level}
— no next-attempt
— replace-result-code code [code...(up to 3 max)]
— no replace-result-code
— peer-address-change-policy {accept | ignore | reject}
— receive-window-size
— no receive-window-size
— rtm-debounce-time debounce-time
— no rtm-debounce-time
— session-limit session-limit
— session-limit unlimited
— no session-limit
— [no] shutdown

configure
— router
— l2tp
— tunnel-selection-blacklist
— add-tunnel never
— add-tunnel on reason>[reason...(up to 8 max)]
— no add-tunnel
— add-tunnel
— max-list-length count
— no max-list-length
— max-time minutes
— no max-time
— timeout-action action
— no timeout-action

```

2.12.1.4 Router Interface Commands

```

config
  — router [router-name]
    — if-attribute
      — admin-group group-name value group-value
      — no admin-group group-name
      — srlg-group group-name value group-value [penalty-weight penalty-weight]
      — no srlg-group group-name
    — [no] interface ip-int-name [unnumbered-mpls-tp | gmpls-loopback | control-
      tunnel]
      — address {ip-address/mask | ip-address netmask} [broadcast {all-ones | host-
        ones}] [track-srrp srrp-instance]
      — no address
      — [no] allow-directed-broadcasts
      — arp-limit limit [log-only] [threshold percent]
      — no arp-limit
      — arp-timeout seconds
      — no arp-timeout
      — bfd transmit-interval [receive receive-interval] [multiplier multiplier] [echo-
        receive echo-interval] [type cpm-np]
      — no bfd
      — cflowd-parameters
      — no cflowd-parameters
        — sampling {unicast | multicast} type {acl | interface} [direction
          {ingress-only | egress-only | both}]
        — no sampling {unicast | multicast}
      — cpu-protection policy-id
      — no cpu-protection
      — description long-description-string
      — no description
      — dhcp
        — description description-string
        — no description
        — gi-address ip-address [src-ip-addr]
        — no gi-address
        — [no] option
          — action {replace | drop | keep}
          — no action
          — circuit-id [ascii-tuple | ifindex | sap-id | vlan-ascii-tuple]
          — no circuit-id
          — remote-id [mac | string string]
          — [no] vendor-specific-option
            — [no] client-mac-address
            — [no] pool-name
            — [no] port-id
            — [no] service-id
            — string text
            — no string
            — [no] system-id
          — python-policy policy-name
          — no python-policy
          — [no] relay-plain-bootp

```

```

— server server1 [server2 ... (up to 8 max)]
— no server
— [no] shutdown
— [no] trusted
— dist-cpu-protection policy-name
— no dist-cpu-protection
— egress
— filter ip ip-filter-id
— filter ipv6 ipv6-filter-id
— no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
— [no] enable-ingress-stats
— [no] enable-mac-accounting
— eth-cfm
— [no] mep mep-id domain md-index association ma-index
— collect-lmm-fc-stats
— fc fc-name [fc-name ... (up to 8 max)]
— no fc
— fc-in-profile fc-name [fc-name ... (up to 8 max)]
— no fc-in-profile
— grace
— eth-ed
— max-rx-defect-window seconds
— no max-rx-defect-window
— priority priority
— no priority
— [no] rx-eth-ed
— [no] tx-eth-ed
— eth-vsm-grace
— [no] rx-eth-vsm-grace
— [no] tx-eth-vsm-grace
— [no] lbm-svc-act-responder
— [no] group-encryption
— encryption-keygroup keygroup-id direction {inbound | outbound}
— no encryption-keygroup direction {inbound | outbound}
— ip-exception filter-id direction {inbound | outbound}
— no ip-exception direction {inbound | outbound}
— hold-time
— up ip seconds
— no up ip
— up ipv6 seconds
— no up ipv6
— down ip seconds [init-only]
— no down
— down ipv6 seconds [init-only]
— no down ipv6
— icmp
— [no] mask-reply
— param-problem [number seconds]
— no param-problem
— redirects [number seconds]
— no redirects
— ttl-expired [number seconds]
— no ttl-expired
— unreachables [number seconds]

```

```

— no unreachable
— if-attribute
— [no] admin-group group-name [group-name...(up to 5 max)]
— no admin-group
— [no] srlg-group group-name [group-name...(up to 5 max)]
— no srlg-group
— ingress
— filter ip ip-filter-id
— filter ipv6 ipv6-filter-id
— no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
— ip-mtu octets
— no ip-mtu
— ip-tunnel
— remote-ip ip-address
— no remote-ip
— lag-link-map-profile lmk-map-profile-id
— no lag-link-map-profile
— lag-per-link-hash class {1 | 2 | 3} weight [1..1024]
— no lag-per-link-hash
— ldp-sync-timer seconds [end-of-lib]
— no ldp-sync-timer
— load-balancing
— egr-ip-load-balancing {source | destination | inner-ip}
— no egr-ip-load-balancing
— lsr-load-balancing hashing-algorithm
— no lsr-load-balancing
— [no] spi-load-balancing
— [no] teid-load-balancing
— [no] local-proxy-arp
— [no] loopback
— mac ieee-mac-addr
— no mac
— network-domain network-domain-name
— no network-domain
— [no] ntp-broadcast
— port port-name
— no port
— [no] proxy-arp-policy
— [no] ptp-hw-assist
— qos-route-lookup [source | destination]
— no qos-route-lookup
— qos network-policy-id [egress-port-redirect-group queue-group-name]
  [egress-instance instance-id] [ingress-fp-redirect-group queue-group-
  name ingress-instance instance-id]
— no qos
— [no] remote-proxy-arp
— secondary {[ip-addr/mask | ip-addr] [netmask]} [broadcast {all-ones | host-
  ones}] [igp-inhibit]
— no secondary [ip-addr/mask | ip-addr] [netmask]
— [no] shutdown
— static-arp ip-addr ieee-mac-addr unnumbered
— no static-arp unnumbered
— [no] strip-label
— tcp-mss mss-value

```

- **no tcp-mss**
- **tos-marking-state** {trusted | untrusted}
- **no tos-marking-state**
- **unnumbered** [ip-addr | ip-int-name]
- **no unnumbered**
- [no] **urpf-check**
 - [no] **ignore-default**
 - **mode** {strict | loose | strict-no-ecmp}
 - **no mode**
- [no] **urpf-selected-vprns**
- **vas-if-type** {to-from-access | to-from-network | to-from-both}
- **no vas-if-type**
- **route-next-hop-policy**
 - [no] **template** template-name
 - **include-group** group-name [pref pref]
 - **no include-group** group-name
 - [no] **exclude-group** group-name
 - [no] **srlg-enable**
 - **protection-type** {link | node}
 - **no protection-type**
 - **nh-type** {ip | tunnel}
 - **no nh-type**

For router interface VRRP commands, see VRRP Configuration Command Reference.

2.12.1.5 Router Interface IPv6 Commands

- ```
config
 — router [router-name]
 — [no] interface ip-int-name
 — [no] ipv6
 — address ipv6-address/prefix-length [eui-64] [preferred] [track-srrp srrp-instance] [modifier cga-modifier]
 — no address ipv6-address/prefix-length
 — bfd transmit-interval [receive receive-interval] [multiplier multiplier] [echo-receive echo-interval [type cpm-np]]
 — no bfd
 — [no] dad-disable
 — icmp6
 — packet-too-big [number seconds]
 — no packet-too-big
 — param-problem [number seconds]
 — no param-problem
 — redirects [number seconds]
 — no redirects
 — time-exceeded [number seconds]
 — no time-exceeded
 — unreachables [number seconds]
 — no unreachables
```

- **link-local-address** *ipv6-address* [**dad-disable**]
- **[no] local-proxy-nd**
- **neighbor** *ipv6-address* [*mac-address*]
- **no neighbor** *ipv6-address*
- **neighbor-limit** *limit* [**log-only**] [*threshold percent*]
- **no neighbor-limit**
- **proxy-nd-policy** *policy-name* [*policy-name...*(up to 5 max)]
- **no proxy-nd-policy**
- **[no] qos-route-lookup**
- **[no] secure-nd**
  - **[no] allow-unsecured-msgs**
  - **link-local-modifier** *modifier*
  - **no link-local-modifier**
  - **public-key-min-bits** *bits*
  - **no public-key-min-bits**
  - **security-parameter** *sec*
  - **no security-parameter**
  - **[no] shutdown**
- **stale-time** *seconds*
- **no stale-time**
- **tcp-mss** *mss-value*
- **no tcp-mss**
- **[no] urpf-check**
  - **mode** {**strict** | **loose** | **strict-no-ecmp**}
  - **no mode**
- **[no] qos-route-lookup**
- **[no] urpf-check**
  - **[no] ignore-default**
  - **mode** {**strict** | **loose**}
  - **no mode**

### 2.12.1.6 Router Advertisement Commands

- ```

config
  — router
    — [no] router-advertisement
      — [no] dns-options
        — servers ipv6-address
        — no servers
        — rdnss-lifetime seconds
        — no rdnss-lifetime
      — [no] interface ip-int-name
        — current-hop-limit number
        — no current-hop-limit
        — [no] dns-options
          — servers ipv6-address
          — no servers
          — rdnss-lifetime {seconds | infinite}
          — no rdnss-lifetime
          — [no] include-dns
        — [no] managed-configuration

```

- **max-advertisement-interval** *seconds*
- **no max-advertisement-interval**
- **min-advertisement-interval** *seconds*
- **no min-advertisement-interval**
- **mtu** *mtu-bytes*
- **no mtu**
- **[no] other-stateful-configuration**
- **prefix** [*ipv6-prefix/prefix-length*]
 - **[no] autonomous**
 - **[no] on-link**
 - **preferred-lifetime** {*seconds* | **infinite**}
 - **no preferred-lifetime**
 - **valid-lifetime** {*seconds* | **infinite**}
 - **no valid-lifetime**
- **reachable-time** *milliseconds*
- **no reachable-time**
- **retransmit-time** *milliseconds*
- **no retransmit-time**
- **router-lifetime** *seconds*
- **no router-lifetime**
- **[no] shutdown**
- **[no] use-virtual-mac**

2.12.2 Command Descriptions

2.12.2.1 Generic Commands

shutdown

| | |
|--------------------|--|
| Syntax | [no] shutdown |
| Context | config>router>if |
| Description | <p>The shutdown command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the no shutdown command.</p> <p>The shutdown command administratively disables an entity. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.</p> <p>Unlike other commands and parameters where the default state is not indicated in the configuration file, shutdown and no shutdown are always indicated in system generated configuration files.</p> <p>The no form of the command puts an entity into the administratively enabled state.</p> |

Default no shutdown

description

| | |
|--------------------|---|
| Syntax | description <i>description-string</i> no description |
| Context | config>router>if>dhcp config>router>if>vrrp config>router>l2tp>group config>router>l2tp>group>tunnel |
| Description | <p>This command creates a text description stored in the configuration file for a configuration context.</p> <p>The no form of the command removes the description string from the context.</p> |
| Default | No description is associated with the configuration context. |
| Parameters | <i>description-string</i> — The description character string. Allowed values are any string up to 80 characters long, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. |

description

| | |
|--------------------|---|
| Syntax | description <i>long-description-string</i> no description |
| Context | config>router>if |
| Description | <p>This command creates a text description stored in the configuration file for a configuration context.</p> <p>The no form of the command removes the description string from the context.</p> |
| Default | No description is associated with the configuration context. |
| Parameters | <i>long-description-string</i> — The description character string. Allowed values are any string up to 160 characters long, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. |

2.12.2.2 Router Global Commands

router

| | | | | | | | | | | |
|--------------------------|---|--|--------------------|--|--|--------------------|--|--|--------------------|-------------------------|
| Syntax | [no] router [router-instance] | | | | | | | | | |
| Context | config | | | | | | | | | |
| Description | <p>This command enables the context to configure router parameters including interfaces, route policies and protocols. This command is also used to create CPM router instances.</p> <p>For CPM router instances, this command enters or creates a user-created CPM router instance. A CPM router instance is not a VPRN router instance. VPRN router instances are configured under configure service vprn. CPM router instances are the only type of non-VPRN router instances that can be created by a user, and they have a user-defined name. CPM router instances only use CPM/CFM/CCM ethernet ports as interfaces.</p> | | | | | | | | | |
| Parameters | <p><i>router-instance</i> — Specifies the router name or CPM router instance.</p> <p>Values</p> <table><tr><td><i>router-instance</i> :</td><td><i>router name</i></td><td></td></tr><tr><td></td><td><i>router-name</i></td><td>Base management <i>cpm-vr-name</i></td></tr><tr><td></td><td><i>cpm-vr-name</i></td><td>[32 characters maximum]</td></tr></table> | <i>router-instance</i> : | <i>router name</i> | | | <i>router-name</i> | Base management <i>cpm-vr-name</i> | | <i>cpm-vr-name</i> | [32 characters maximum] |
| <i>router-instance</i> : | <i>router name</i> | | | | | | | | | |
| | <i>router-name</i> | Base management <i>cpm-vr-name</i> | | | | | | | | |
| | <i>cpm-vr-name</i> | [32 characters maximum] | | | | | | | | |
| Default | Base | | | | | | | | | |

aggregate

| | |
|--------------------|---|
| Syntax | aggregate <i>ip-prefix/ip-prefix-length</i> [summary-only] [as-set] [aggregator <i>as-number:ip-address</i>] [black-hole] [community <i>comm-id</i>] [description <i>description</i>] aggregate <i>ip-prefix/ip-prefix-length</i> [summary-only] [as-set] [aggregator <i>as-number:ip-address</i>] [community <i>comm-id</i>] [indirect <i>ip-address</i>] [description <i>description</i>] no aggregate <i>ip-prefix/ip-prefix-length</i> |
| Context | config>router |
| Description | <p>This command creates an aggregate route.</p> <p>Use this command to automatically install an aggregate in the routing table when there are one or more component routes. A component route is any route used for forwarding that is a more-specific match of the aggregate.</p> <p>The use of aggregate routes can reduce the number of routes that need to be advertised to neighbor routers, leading to smaller routing table sizes.</p> |

Overlapping aggregate routes may be configured; in this case a route becomes a component of only the one aggregate route with the longest prefix match. For example if one aggregate is configured as 10.0.0.0/16 and another as 10.0.0.0/24, then route 10.0.128/17 would be aggregated into 10.0.0.0/16, and route 10.0.0.128/25 would be aggregated into 10.0.0.0/24. If multiple entries are made with the same prefix and the same mask the previous entry is overwritten.

A standard 4-byte BGP community may be associated with an aggregate route in order to facilitate route policy matching.

By default aggregate routes are not installed in the forwarding table, however there are configuration options that allow an aggregate route to be installed with a black-hole next hop or with an indirect IP address as next hop.

The **no** form of the command removes the aggregate.

Default No aggregate routes are defined.

Parameters *ip-prefix* — The destination address of the aggregate route in dotted decimal notation.

Values The following values apply to the 7750 SR and 7950 XRS:

| | | | |
|--------------------|-------------------------------------|--------------|--|
| ipv4-prefix | a.b.c.d (host bits must be 0) | | |
| ipv4-prefix-length | 0 to 32 | | |
| ipv6-prefix | x:x:x:x:x:x:x (eight 16-bit pieces) | | |
| | x:x:x:x:x:d.d.d.d | | |
| | x: | [0 to FFFF]H | |
| | d: | [0 to 255]D | |
| ipv6-prefix-length | 0 to 128 | | |

Values The following values apply to the 7450 ESS:

| | |
|--------------------|-------------------------------|
| ipv4-prefix | a.b.c.d (host bits must be 0) |
| ipv4-prefix-length | 0 to 32 |

ip-prefix-length — The mask associated with the network address expressed as a mask length.

Values 0 to 32

summary-only — This optional parameter suppresses advertisement of more specific component routes for the aggregate.

To remove the **summary-only** option, enter the same aggregate command without the **summary-only** parameter.

as-set — This optional parameter is only applicable to BGP and creates an aggregate where the path advertised for this route will be an AS_SET consisting of all elements contained in all paths that are being summarized. Use this feature carefully as it can increase the amount of route churn due to best path changes.

aggregator as-number:ip-address — This optional parameter specifies the BGP aggregator path attribute to the aggregate route. When configuring the aggregator, a two-octet AS number used to form the aggregate route must be entered, followed by the IP address of the BGP system that created the aggregate route.

community — This configuration option associates a BGP community with the aggregate route. The community can be matched in route policies and is automatically added to BGP routes exported from the aggregate route.

comm-id — Specifies community IDs, up to 72 characters.

Values [2 byte-as-number:comm-val | well-known-comm]

where:

- 2 byte-as-number — 0 to 65535
- comm-val — 0 to 65535
- well-known-comm — **no-export** | **no-export-subconfed** | **no-advertise**

black-hole — This optional parameter installs the aggregate route, when activated, in the FIB with a black-hole next-hop; where packets matching this route are discarded.

indirect ip-address — This configuration option specifies that the aggregate route should be installed in the FIB with a next-hop taken from the route used to forward packets to ip-address.

Values The following values apply to the 7750 SR and 7950 XRS:

| | |
|-------------|---------------------|
| ipv4-prefix | a.b.c.d |
| ipv6-prefix | x:x:x:x:x:x:x |
| | x:x:x:x:x:x:d.d.d.d |
| | x: [0 to FFFF]H |
| | d: [0 to 255]D |

Values The following values apply to the 7450 ESS:

ipv4-prefix: a.b.c.d

description description-text — Specifies a text description stored in the configuration file for a configuration context.

autonomous-system

Syntax **autonomous-system** *autonomous-system*
no autonomous-system

| | |
|--------------------|---|
| Context | config>router |
| Description | <p>This command configures the autonomous system (AS) number for the router. A router can only belong to one AS. An AS number is a globally unique number with an AS. This number is used to exchange exterior routing information with neighboring ASs and as an identifier of the AS itself.</p> <p>If the AS number is changed on a router with an active BGP instance, the new AS number is not used until the BGP instance is restarted either by administratively disabling/enabling (shutdown/no shutdown) the BGP instance or rebooting the system with the new configuration.</p> |
| Default | No autonomous system number is defined. |
| Parameters | <i>autonomous-system</i> — The autonomous system number expressed as a decimal integer. |
| Values | 1 to 4294967295 |

class-forwarding

| | |
|--------------------|---|
| Syntax | [no] class-forwarding |
| Context | config>router |
| Description | <p>This command enables class-based forwarding (CBF) over IGP shortcuts. When the class-forwarding command is enabled, the following types of packets are forwarded based on their forwarding class:</p> <ul style="list-style-type: none">• packets of BGP prefixes• CPM originated packets for the families (IPv4 only, IPv6 only, or both IPv4 and IPv6) which have been enabled over IGP shortcuts using the igp-shortcut CLI context in one or more IGP instances <p>The SR OS CBF implementation supports spraying of packets over a maximum of four forwarding sets of ECMP LSPs. The user must define a class-forwarding policy object in MPLS to configure the mapping of FCs to the forwarding sets. Then, the user assigns the CBF policy name and set ID to each MPLS LSP that is used in IGP shortcuts.</p> <p>When a BGP IPv4 or IPv6 prefix is resolved, the FC of the packet is used to look up the forwarding set ID. Then, a modulo operation is performed on the tunnel next-hops of this set ID only, to spray packets of this FC. The data path concurrently implements CBF and ECMP within the tunnels of each set ID.</p> <p>CPM-originated packets on the router, including control plane and OAM packets, are forwarded over a single LSP from the set of LSPs that the packet's FC is mapped to, as per the CBF configuration.</p> |



Note: Weighted ECMP, at the transport tunnel level of BGP prefixes over IGP shortcuts and the CBF feature on a per BGP next-hop basis are mutually exclusive.

Default no class-forwarding

confederation

Syntax **confederation** *confed-as-num* **members** *as-number* [*as-number...up to 15 max*]
no confederation [*confed-as-num members as-number...up to 15 max*]

Context config>router

Description This command creates confederation autonomous systems within an AS.

This technique is used to reduce the number of IBGP sessions required within an AS. Route reflection is another technique that is commonly deployed to reduce the number of IBGP sessions.

The **no** form of the command deletes the specified member AS from the confederation.

When no members are specified in the **no** statement, the entire list is removed and **confederation** is disabled.

When the last member of the list is removed, **confederation** is disabled.

Default no confederation - no confederations are defined.

Parameters *confed-as-num* — The confederation AS number expressed as a decimal integer.

Values 1 to 65535

members *member-as-num* — The AS number of members that are part of the confederation, expressed as a decimal integer. Up to 15 members per *confed-as-num* can be configured.

Values 1 to 65535

ecmp

Syntax **ecmp** *max-ecmp-routes*
no ecmp

Context config>router

Description This command enables ECMP and configures the number of routes for path sharing; for example, the value 2 means two equal cost routes will be used for cost sharing.

ECMP can only be used for routes learned with the same preference and same protocol.

When more ECMP routes are available at the best preference than configured in *max-ecmp-routes*, then the lowest next-hop IP address algorithm is used to select the number of routes configured in *max-ecmp-routes*.

The **no** form of the command disables ECMP path sharing. If ECMP is disabled and multiple routes are available at the best preference and equal cost, then the route with the lowest next-hop IP address is used.

Default no ecmp

Parameters *max-ecmp-routes* — The maximum number of equal cost routes allowed on this routing table instance, expressed as a decimal integer. Setting ECMP *max-ecmp-routes* to 1 yields the same result as entering **no ecmp**.

Values 0 to 32

entropy-label

Syntax **entropy-label**
no entropy-label

Context config>router

Description If **entropy-label** is configured, the Entropy label and Entropy Label Indicator is inserted on packets for which at least one LSP in the stack for the far-end of the LDP or RSVP tunnel used by an IGP or BGP shortcut has advertised entropy-label-capability. If the tunnel is of type RSVP, then **entropy-label** must also have been enabled under **config>router>mpls** or **config>router>mpls>lsp**.

This configuration will result in other traffic that is forwarded over an LDP or RSVP LSP for which this router is the LER, and for which there is no explicit service endpoint on this router, to have the EL/ELI enabled, subject to the LSP far-end advertising entropy-label-capability. An example of such traffic includes packets arriving on a stitched LDP LSP forwarded over an RSVP LSP.

Default no entropy-label

flowspec

Syntax **flowspec**

Context config>router

Description This command enables the context to configure flowspec-related parameters for the specified routing instance.

ip-filter-max-size

| | |
|--------------------|--|
| Syntax | ip-filter-max-size { <i>value</i> default } |
| Context | config>router>flowspec |
| Description | <p>This command configures the maximum number of flowspec routes or rules that can be embedded into the auto-created embedded filter (fSpec-X). Flowspec filter entries embedded in a filter policy in this routing instance will use filter entries from the range between “embedding offset + 1” and “embedding offset + ip-filter-max-size”.</p> <p>The sum of the ip-filter-max-size <i>value</i> parameter and the highest offset in any IPv4 filter that embeds IPv4 flowspec rules from this routing instance (excluding filters that embed at offset 65535) must not exceed 65535.</p> <p>The ip-filter-max-size configuration can be adjusted up or down at any time. If the number of IPv4 flowspec rules that are currently installed is <i>M</i>, and the new limit is <i>N</i>, where $N < M$, then the last set of rules from <i>N</i> to <i>M</i> (by flowspec order) are immediately removed, but are retained in the BGP RIB. If the limit is increased, new rules are programmed only as they are received again in new BGP updates.</p> |
| Default | ip-filter-max-size default |
| Parameters | <p><i>value</i> — The maximum number of flowspec routes or rules that can be embedded into an ingress IP filter policy.</p> <p>Values 0 to 65535</p> <p>default — Keyword to configure the maximum size as 512.</p> |

ipv6-filter-max-size

| | |
|--------------------|---|
| Syntax | ipv6-filter-max-size { <i>value</i> default } |
| Context | config>router>flowspec |
| Description | <p>This command configures the maximum number of IPv6 flowspec routes or rules that can be embedded into the auto-created embedded filter (fSpec-X). Flowspec filter entries embedded in a filter policy in this routing instance will use filter entries from the range between “embedding offset + 1” and “embedding offset + ip-filter-max-size”.</p> <p>The sum of the ip-filter-max-size <i>value</i> parameter and the highest offset in any IPv6 filter that embeds IPv6 flowspec rules from this routing instance (excluding filters that embed at offset 65535) must not exceed 65535.</p> <p>The ip-filter-max-size configuration can be adjusted up or down at any time. If the number of IPv6 flowspec rules that are currently installed is <i>M</i>, and the new limit is <i>N</i>, where $N < M$, then the last set of rules from <i>N</i> to <i>M</i> (by flowspec order) are immediately removed, but are retained in the BGP RIB. If the limit is increased, new rules are programmed only as they are received again in new BGP updates.</p> |

| | |
|-------------------|--|
| Default | ipv6-filter-max-size default |
| Parameters | <i>value</i> — The maximum number of flowspec routes or rules that can be embedded into an ingress IP filter policy. Values 0 to 65535 default — Keyword to configure the maximum size as 512. |

weighted-ecmp

| | |
|--------------------|---|
| Syntax | [no] weighted-ecmp |
| Context | config>router |
| Description | <p>This command enables the weighted load-balancing, or weighted ECMP, over MPLS LSP.</p> <p>When this command is enabled, packets of IGP, BGP, and static route prefixes resolved to a set of ECMP tunnel next-hops are sprayed proportionally to the weights configured for each MPLS LSP in the ECMP set.</p> <p>Weighted load-balancing over MPLS LSP is supported in the following forwarding contexts:</p> <ul style="list-style-type: none">• IGP prefix resolved to IGP shortcuts in RTM (igp-shortcut or advertise-tunnel-link enabled in the IGP instance).• BGP prefix with the BGP next-hop resolved to IGP shortcuts in RTM (igp-shortcut or advertise-tunnel-link enabled in the IGP instance).• Static route prefix resolved to an indirect next-hop, which itself is resolved to a set of equal-metric MPLS LSPs in TTM. The user can allow automatic selection or specify the names of the equal-metric MPLS LSPs in TTM to be used in the ECMP set.• Static route prefix resolved to an indirect next-hop, which is resolved to IGP shortcuts in RTM.• BGP prefix with a BGP next-hop resolved to a static route, which resolves to a set of tunnel next-hops toward an indirect next-hop in RTM or TTM.• BGP prefix resolving to another BGP prefix, whose next-hop is resolved to a set of ECMP tunnel next-hops with a static route in RTM or TTM or to IGP shortcuts in RTM. |

IGP computes the normalized weight for each prefix tunnel next-hop. IGP updates the route in RTM with the set of tunnel next-hops and normalized weights. RTM downloads the information to IOM for inclusion in the FIB.

If one or more LSPs in the ECMP set of a prefix do not have a weight configured, the regular ECMP spraying for the prefix will be performed.

The weight assigned to an LSP impacts only the forwarding decision, not the routing decision. In other words, it does not change the selection of the set of ECMP tunnel next-hops of a prefix when more next-hops exist than the value of the router **ecmp** option. Once the set of tunnel next-hops is selected, the LSP weight is used to modulate the amount of packets forwarded over each next-hop. It also does not change the hash routine, but only the spraying of the flows over the tunnel next-hops is modified to reflect the normalized weight of each tunnel next-hop.

The no version of the command resumes regular ECMP spraying of packets of IGP, BGP, and static route prefixes over MPLS LSP.

Default no weighted-ecmp

fib-priority

Syntax fib-priority {high | standard}

Context config>router

Description This command specifies the FIB priority for VPRN.

Default fib-priority standard

icmp-tunneling

Syntax icmp-tunneling
no icmp-tunneling

Context config>router

Description This command enables the tunneling of ICMP reply packets over MPLS LSP at a LSR node as per RFC 3032.

The LSR part of this feature consists of crafting the reply ICMP packet of type=11- 'time exceeded', with a source address set to a local address of the LSR node, and appending the IP header and leading payload octets of the original datagram. The system skips the lookup of the source address of the sender of the label TTL expiry packet, which becomes the destination address of the ICMP reply packet. Instead, CPM injects the ICMP reply packet in the forward direction of the MPLS LSP the label TTL expiry packet was received from. The TTL of pushed labels should be set to 255.

The source address of the ICMP reply packet is determined as follows. The LSR uses the address of the outgoing interface for the MPLS LSP. With LDP LSP or BGP LSP multiple ECMP next-hops can exist and in such a case the first outgoing interface is selected. If that interface does not have an address of the same family (IPv4 or IPv6) as the ICMP packet, then the system address of the same family is selected. If one is not configured, the packet is dropped.

When the packet is received by the egress LER, it performs a regular user packet lookup in the data path in the GRT context for BGP shortcut, 6PE, and BGP label route prefixes, or in VPRN context for VPRN and 6VPE prefixes. It then forwards it to the destination, which is the sender of the original packet which TTL expired at the LSR.

If the egress LER does not have a route to the destination of the ICMP packet, it drops the packets.

The rate of the tunneled ICMP replies at the LSR can be directly or indirectly controlled by the existing IOM level and CPM levels mechanisms. Specifically, the rate of the incoming UDP traceroute packets received with a label stack can be controlled at ingress IOM using the distributed CPU protection feature. The rate of the ICMP replies by CPM can also be directly controlled by configuring a system wide rate limit for packets ICMP replies to MPLS expired packets which are successfully forwarded to CPM using the command 'configure system security vprn-network-exceptions'. While this command's name refers to VPRN service, this feature rate limits ICMP replies for packets received with any label stack, including VPRN and shortcuts.

The 7450 ESS, 7750 SR, and 7950 XRS implementation supports appending to the ICMP reply of type Time Exceeded the MPLS label stack object defined in RFC 4950. It does not include it in the ICMP reply type of Destination unreachable.

The new MPLS Label Stack object permits an LSR to include label stack information including label value, EXP, and TTL field values, from the encapsulation header of the packet that expired at the LSR node. The ICMP message continues to include the IP header and leading payload octets of the original datagram.

In order to include the MPLS Label Stack object, SR OS implementation adds support of RFC 4884 which defines extensions for a multi-part ICMPv4/v6 message of type Time Exceeded.

The **no** form of command disables the tunneling of ICMP reply packets over MPLS LSP at a LSR node.

Default no icmp-tunneling

ip-fast-reroute

Syntax [no] ip-fast-reroute

Context config>router

Description This command enables IP Fast-Reroute (FRR) feature on the system.

This feature provides for the use of a Loop-Free Alternate (LFA) backup next-hop for forwarding in-transit and CPM generated IP packets when the primary next-hop is not available. IP FRR is supported on IPv4 and IPv6 OSPF/IS-IS prefixes forwarded in the base router instance to a network IP interface or to an IES SAP interface or spoke interface. It is also supported for VPRN VPN-IPv4 OSPF prefixes and VPN-IPv6 OSPF prefixes forwarded to a VPRN SAP interface or spoke interface.

IP FRR also provides a LFA backup next-hop for the destination prefix of a GRE tunnel used in an SDP or in VPRN auto-bind.

When any of the following events occurs, IGP instructs in the fast path on the XMA to enable the LFA backup next-hop:

- OSPF/IS-IS interface goes operationally down: physical or local admin shutdown.
- Timeout of a BFD session to a next-hop when BFD is enabled on the OSPF/IS-IS interface

When the SPF computation determines there is more than one primary next-hop for a prefix, it will not program any LFA next-hop in RTM. Therefore, the IP prefix will resolve to the multiple equal-cost primary next-hops that provide the required protection.

The **no** form of this command disables the IP FRR feature on the system

Default no ip-fast-reroute

mc-maximum-routes

Syntax **mc-maximum-routes** *number* [**log-only**] [**threshold** *threshold*]
no mc-maximum-routes

Context config>router

Description This command specifies the maximum number of multicast routes that can be held within a VPN routing/forwarding (VRF) context. When this limit is reached, a log and SNMP trap are sent. If the **log-only** parameter is not specified and the maximum-routes value is set below the existing number of routes in a VRF, then no new joins will be processed.

The **no** form of the command disables the limit of multicast routes within a VRF context. Issue the **no** form of the command only when the VPRN instance is shutdown.

Default no mc-maximum-routes

Parameters *number* — Specifies the maximum number of routes to be held in a VRF context

Values 1 to 2147483647

log-only — Specifies that if the maximum limit is reached, only log the event. **log-only** does not disable the learning of new routes

threshold — The percentage at which a warning log message and SNMP trap should be sent.

Values 0 to 100

Default 10

mpls-labels

Syntax **mpls-labels**

Context config>router

Description This command creates a context for the configuration of global parameters related to MPLS labels.

static-label-range

Syntax **static-label-range** *static-range*
no static-label-range

Context config>router>mpls-labels

Description This command configures the range of MPLS static label values shared among static LSP, MPLS-TP LSP, and static service VC label. Once this range is configured, it is reserved and cannot be used by other protocols such as RSVP, LDP, BGP, or Segment Routing to assign a label dynamically.

Default static-label-range 18400

Parameters *static-range* — Specifies the size of the static label range in number of labels. The minimum label value in the range is 32. The maximum label value is therefore computed as {32+ static-range-1}.

Values 0 to 262112

Default 18400

bgp-labels-hold-timer

Syntax **bgp-labels-hold-timer** *seconds*
[no] bgp-labels-hold-timer

Context config>router>mpls-labels

Description This command configures the BGP labels hold timer on the ingress router.

Default 0

Parameters *seconds* — Specifies the seconds
Values 0 to 255

sr-labels

Syntax **sr-labels start** *start-value* **end** *end-value*
no sr-labels

Context config>router>mpls-labels

Description This command configures the range of the Segment Routing Global Block (SRGB). It is a label block which is used for assigning labels to segment routing prefix SIDs originated by this router. This range is carved from the system dynamic label range and is not instantiated by default.

 This is a reserved label and once configured it cannot be used by other protocols such as RSVP, LDP, and BGP to assign a label dynamically.

Default no sr-labels

Parameters *start-value* — Specifies the start label value in the SRGB
Values 18432 to 524287 within dynamic label range
end-value — Specifies the end label value in the SRGB
Values 18432 to 524287 within dynamic label range

mss-adjust-group

Syntax **mss-adjust-group** *nat-group-id* **segment-size** *segment-size*
no mss-adjust-group

Context config>router
 config>service>vprn

Description This command associates the MSS adjust group consisting of multiple ISAs with the routing context in which the application requiring TCP MSS adjust resides.

Parameters *nat-group-id* — Specifies the NAT group used for TCP MSS adjust
 segment-size — Specifies the value to put into the TCP Maximum Segment Size (MSS) option if it is not already present, or if the present value is higher

multicast-info-policy

Syntax **multicast-info-policy** *policy-name*

no multicast-info-policy

| | |
|--------------------|---|
| Context | config>router |
| Description | This command configures multicast information policy. |
| Default | no multicast-info-policy |
| Parameters | <i>policy-name</i> — Specifies the policy name |
| Values | 32 chars max |

network-domains

| | |
|--------------------|---|
| Syntax | network-domains |
| Context | config>router |
| Description | This command opens context for defining network-domains. This command is applicable only in the base routing context. |

description

| | |
|--------------------|--|
| Syntax | [no] description <i>string</i> |
| Context | config>router>network-domains>network-domain |
| Description | This command creates a text description stored in the configuration file for a configuration context. The no form of the command removes the description string from the context. |
| Default | no description |
| Parameters | <i>string</i> — Specifies the description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special character (#, \$, space, and so on), the entire string must be enclosed within double quotes. |

network-domain

| | |
|--------------------|--|
| Syntax | network-domain <i>network-domain-name</i> [create] no network-domain <i>network-domain-name</i> |
| Context | config>router>network-domains |
| Description | This command creates network-domains that can be associated with individual interfaces and SDPs. |

| | |
|-------------------|--|
| Default | network-domain "default" |
| Parameters | <i>network-domain-name</i> — Network domain name character string. |

rpki-session

| | |
|--------------------|---|
| Syntax | rpki-session <i>ip-address</i> no rpki-session <i>ip-address</i> |
| Context | config>router>origin-validation |
| Description | This command configures a session with an RPKI local cache server by using the RPKI-Router protocol. It is over these sessions that the router learns dynamic VRP entries expressing valid origin AS and prefix associations. SR OS supports the RPKI-Router protocol over TCP/IPv4 or TCP/IPv6 transport. The router can set up an RPKI-Router session using the base routing table (in-band) or the management router (out-of-band). Configure the command in the config>router management instance to configure a session using the management port. |
| Default | no rpki-session |
| Parameters | <i>ip-address</i> — An IPv4 address or an IPv6 address. If the IPv6 address is link-local then the interface name must be appended to the IPv6 address after a hyphen (-). |

connect-retry

| | |
|--------------------|---|
| Syntax | connect-retry <i>seconds</i> no connect-retry |
| Context | config>router>origin-validation>rpki-session |
| Description | This command configures the time in seconds to wait between one TCP connection attempt that fails and the next attempt. The default (with no connect-retry) is 120 seconds. |
| Default | no connect-retry |
| Parameters | <i>seconds</i> — Specifies time in seconds. Values 1 to 65535 |

description

| | |
|----------------|---|
| Syntax | description <i>description-string</i> no description |
| Context | config>router>origin-validation>rpki-session |

| | |
|--------------------|--|
| Description | This command configures a description for an RPKI-Router session. |
| Default | no description |
| Parameters | <i>description-string</i> — Specifies a text string up to 80 characters in length. |

local-address

| | |
|--------------------|---|
| Syntax | local-address <i>ip-address</i> no local-address |
| Context | config>router>origin-validation>rpki-session |
| Description | This command configures the local address to use for setting up the TCP connection used by an RPKI-Router session. The default local-address is the outgoing interface IPv4 or IPv6 address. The local-address cannot be changed without first shutting down the session. |
| Default | no local-address |
| Parameters | <i>ip-address</i> — Specifies an IPv4 address or an IPv6 address. |

port

| | |
|--------------------|--|
| Syntax | port <i>port-id</i> no port |
| Context | config>router>origin-validation>rpki-session |
| Description | This command configures the destination port number to use when contacting the cache server. The default port number is 323. The port cannot be changed without first shutting down the session. |
| Default | no port |
| Parameters | <i>port-id</i> — Specifies a port-id. Values 0 to 65535 |

refresh-time

| | |
|----------------|--|
| Syntax | refresh-time <i>seconds1</i> hold-time <i>seconds2</i> no refresh-time |
| Context | config>router>origin-validation>rpki-session |

| | |
|--------------------|--|
| Description | <p>This command is used to configure the refresh-time and hold-time intervals that are used for liveness detection of the RPKI-Router session. The refresh-time defaults to 300 seconds and is reset whenever a Reset Query PDU or Serial Query PDU is sent to the cache server. When the timer expires, a new Serial Query PDU is sent with the last known serial number.</p> <p>The hold-time specifies the length of time in seconds that the session is to be considered UP without any indication that the cache server is alive and reachable. The timer defaults to 600 seconds and must be at least 2x the refresh-time (otherwise the CLI command is not accepted). Reception of any PDU from the cache server resets the hold timer. When the hold-time expires, the session is considered to be DOWN and the stale timer is started.</p> |
| Default | no refresh-time |
| Parameters | <p><i>seconds1</i> — Specifies a time in seconds.</p> <p>Values 30 to 32767</p> <p><i>seconds2</i> — Specifies a time in seconds.</p> <p>Values 60 to 65535</p> |

shutdown

| | |
|--------------------|---|
| Syntax | shutdown no shutdown |
| Context | config>router>origin-validation>rpk-session |
| Description | This command administratively disables an RPKI-Router session. The no form of the command enables the RPKI-Router session. |
| Default | no shutdown |

stale-time

| | |
|--------------------|--|
| Syntax | stale-time seconds no stale-time |
| Context | config>router>origin-validation>rpk-session |
| Description | This command configures the maximum length of time that prefix origin validation records learned from the cache server remain usable after the RPKI-Router session goes down. The default stale-time is 3600 seconds (1 hour). When the timer expires all remaining stale entries associated with the session are deleted. |
| Default | no stale-time |
| Parameters | <p><i>seconds</i> — Specifies a time in seconds.</p> <p>Values 60 to 3600</p> |

static-entry

| | |
|--------------------|---|
| Syntax | static-entry <i>ip-prefix/ip-prefix-length upto prefix-length2 origin-as as-number</i> [valid invalid] no static-entry <i>ip-prefix/ip-prefix-length upto prefix-length2 origin-as as-number</i> |
| Context | config>router>origin-validation |
| Description | <p>This command configures a static VRP entry indicating that a specific origin AS is either valid or invalid for a specific IP prefix range. Static VRP entries are stored along with dynamic VRP entries (learned from local cache servers using the RPKI-Router protocol) in the origin validation database of the router. This database is used for determining the origin-validation state of IPv4 and/or IPv6 BGP routes received over sessions with the enable-origin-validation command configured.</p> <p>Static entries can only be configured under the config>router>origin-validation context of the base router.</p> |
| Default | no static entries |
| Parameters | <p><i>ip-prefix/ip-prefix-length</i> — Specifies an IPv4 or IPv6 address with a minimum prefix length value.</p> <p>Values 60 to 3600</p> <p><i>prefix-length2</i> — Specifies the maximum prefix length.</p> <p><i>as-number</i> — Specifies as-number.</p> <p>Values 0 to 4294967295</p> <p>valid — Specifies a keyword meaning the static entry expresses a valid combination of origin AS and prefix range.</p> <p>invalid — Specifies a keyword meaning the static entry expresses an invalid combination of origin AS and prefix range.</p> |

router-id

| | |
|--------------------|--|
| Syntax | router-id <i>ip-address</i> no router-id |
| Context | config>router |
| Description | <p>This command configures the router ID for the router instance.</p> <p>The router ID is used by both OSPF and BGP routing protocols in this instance of the routing table manager. IS-IS uses the router ID value as its system ID.</p> <p>When configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time a protocol is initialized, the new router ID is used. This can result in an interim period of time when different protocols use different router IDs.</p> |

It is possible to configure SR OS to operate with an IPv6 only BOF and no IPv4 system interface address. When configured in this manner, the operator must explicitly define IPv4 router IDs for protocols such as OSPF and BGP as there is no mechanism to derive the router ID from an IPv6 system interface address.

To force the new router ID to be used, issue the **shutdown** and **no shutdown** commands for each protocol that uses the router ID, or restart the entire router.

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

Default The system uses the system interface address (which is also the loopback address).

If a system interface address is not configured, use the last 32 bits of the chassis MAC address.

Parameters *router-id* — The 32 bit router ID expressed in dotted decimal notation or as a decimal value.

service-prefix

Syntax **service-prefix** *ip-prefix/mask* | *ip-prefix netmask* [**exclusive**]
no service-prefix *ip-prefix/mask* | *ip-prefix netmask*

Context config>router

Description This command creates an IP address range reserved for IES or VPLS services.

The purpose of reserving IP addresses using **service-prefix** is to provide a mechanism to reserve one or more address ranges for services.

When services are defined, the address must be in the range specified as a service prefix. If a service prefix is defined, then IP addresses assigned for services must be within one of the ranges defined in the **service-prefix** command. If the **service-prefix** command is not configured, then no limitations exist.

Addresses in the range of a service prefix can be allocated to a network port unless the exclusive parameter is used. Then, the address range is exclusively reserved for services.

When a range that is a superset of a previously defined service prefix is defined, the subset is replaced with the superset definition; for example, if a service prefix exists for 10.10.10.0/24, and a service prefix is configured as 10.10.0.0/16, then 10.10.10.0/24 is replaced by the new 10.10.0.0/16 configuration.

When a range that is a subset of a previously defined service prefix is defined, the subset replaces the existing superset, providing addresses used by services are not affected; for example, if a service prefix exists for 10.10.0.0/16, and a service prefix is configured as 10.10.10.0/24, then the 10.10.0.0/16 entry is removed as long as no services are configured that use 10.10.x.x addresses other than 10.10.10.x.

The **no** form of the command removes all address reservations. A service prefix cannot be removed while one or more service uses an address or addresses in the range.

Default no service-prefix - No IP addresses are reserved for services.

Parameters *ip-prefix/mask* — The IP address prefix to include in the service prefix allocation in dotted decimal notation.

Values

| | | |
|---------------------|-------------------------------------|--------------|
| ipv4-prefix: | a.b.c.d (host bits must be 0) | |
| ipv4-prefix-length: | 0 to 32 | |
| ipv6-prefix: | x:x:x:x:x:x:x (eight 16-bit pieces) | |
| | x:x:x:x:x:x.d.d.d.d | |
| | x: | [0 to FFFF]H |
| | d: | [0 to 255]D |
| ipv6-prefix-length: | 0 to 128 | |

Values

exclusive

When this option is specified, the addresses configured are exclusively used for services and cannot be assigned to network ports.

sgt-qos

Syntax **sgt-qos**

Context config>router

Description This command configures DSCP/dot1p re-marking for self-generated traffic.

application

Syntax **application** *dscp-app-name* **dscp** {*dscp-value* | *dscp-name*}
application *dot1p-app-name* **dot1p** *dot1p-priority*
no application {*dscp-app-name* | *dot1p-app-name*}

Context config>router>sgt-qos

Description This command configures DSCP/dot1p remarking for self-generated traffic. When an application is configured using this command, then the specified DSCP name/value is used for all packets generated by this application.

Using the value configured in this command:

- sets the DSCP bits in the IP packet
- maps to the FC. This value will be signaled from the CPM to the egress forwarding complex.
- based on this signaled FC, the egress forwarding complex QoS policy sets the IEEE 802.1p and MPLS EXP bits
- the DSCP value in the egress IP header will be as configured in this command. The egress QoS policy will not overwrite this value.

Only one DSCP name/value can be configured per application, if multiple entries are configured then the subsequent entry overrides the previous configured entry.

The **no** form of this command resets the command back to the default value.

| | |
|-------------------|---|
| Parameters | <p><i>dscp-app-name</i> — Specifies the DSCP application name.</p> <p>Values bgp, cflowd, dhcp, diameter, dns, ftp, gtp, icmp, igmp, igmp-reporter, l2tp, ldp, mld, msdp, ndis, ntp, ospf, pcep, pim, ptp, radius, rip, rsvp, sflow, snmp, snmp-notification, srtp, ssh, syslog, tacplus, telnet, tftp, traceroute, vrrp</p> <p><i>dscp-value</i> — Specifies the DSCP value.</p> <p>Values 0 to 63</p> <p><i>dscp-name</i> — Specifies the DSCP name.</p> <p>none, be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63</p> <p><i>dot1p-priority</i> — Specifies the dot1p priority.</p> <p>Values none, 0 to 7</p> <p><i>dot1p-app-name</i> — Specifies the dot1p application name.</p> <p>Values arp, isis, pppoe</p> |
|-------------------|---|

dscp

| | |
|--------------------|---|
| Syntax | <p>dscp <i>dscp-name</i> fc <i>fc-name</i></p> <p>no dscp <i>dscp-name</i></p> |
| Context | config>router>sgt-qos |
| Description | This command creates a mapping between the DiffServ Code Point (DSCP) of the self-generated traffic and the forwarding class. |

Self-generated traffic that matches the specified DSCP will be assigned to the corresponding forwarding class. Multiple commands can be entered to define the association of some or all sixty-four DiffServ code points to the forwarding class. For undefined code points, packets are assigned to the forwarding class specified under the default-action command.

All DSCP names that defines a DSCP value must be explicitly defined.

The **no** form of this command removes the DiffServ code point to forwarding class association. The default-action then applies to that code point value.

| | |
|-------------------|---|
| Parameters | <i>dscp-name</i> — Specifies the DSCP name. |
| Values | be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63 |
| | <i>fc-name</i> — Specifies the forwarding class name. |
| Values | be, l2, af, l1, h2, ef, h1, nc |

bfd-template

| | |
|--------------------|---|
| Syntax | bfd-template <i>name</i> [create] no bfd-template <i>name</i> |
| Context | config>router>bfd |
| Description | This command creates or edits a BFD template. A BFD template defines the set of configurable parameters used by a BFD session. These include the transmit and receive timers used for BFD CC packets, the transmit timer interval used when the session is providing a CV function, the multiplier value, the echo-receive interval, and whether the BFD session terminates in the CPM network processor. |
| Default | no bfd-template |
| Parameters | <i>name</i> — Specifies a text string name for the template up to 32 characters in printable 7-bit ASCII, enclosed in double quotes. |

transmit-interval

| | |
|----------------|--|
| Syntax | transmit-interval <i>transmit-interval</i> no transmit-interval |
| Context | config>router>bfd>bfd-template |

| | |
|--------------------|--|
| Description | This command specifies the transmit timer used for BFD packets. If the template is used for a BFD session on an MPLS-TP LSP, then this timer is used for CC packets. |
| Default | transmit-interval 100 |
| Parameters | <i>transmit-interval</i> — Specifies the transmit interval. The minimum interval that can be configured is hardware dependent. |
| Values | 10 ms to 100,000 ms in 1 ms intervals |
| Default | 10 ms for CPM3 or higher; 1 second for other hardware |

receive-interval

| | |
|--------------------|---|
| Syntax | receive-interval <i>receive-interval</i> no receive-interval |
| Context | config>router>bfd>bfd-template |
| Description | This command specifies the receive timer used for BFD packets. If the template is used for a BFD session on an MPLS-TP LSP, then this timer is used for CC packets. |
| Default | receive-interval 100 |
| Parameters | <i>receive-interval</i> — Specifies the receive interval. The minimum interval that can be configured is hardware dependent. |
| Values | 10 ms to 100,000 ms in 1 ms intervals |
| Default | 10 ms for CPM3 or higher; 1 second for other hardware |

echo-receive

| | |
|--------------------|---|
| Syntax | echo-receive <i>echo-interval</i> no echo-receive |
| Context | config>router>bfd>bfd-template |
| Description | This command sets the minimum echo receive interval, in milliseconds, for a session. This is not used by a BFD session for MPLS-TP. |
| Default | echo-receive 100 |
| Parameters | <i>echo-interval</i> — Specifies the echo receive interval. |
| Values | 100 ms to 100,000 ms in 1 ms increments |
| Default | 100 |

multiplier

| | | | | | |
|--------------------|---|---------------|-------------------|----------------|---|
| Syntax | multiplier <i>multiplier</i> no multiplier | | | | |
| Context | config>router>bfd>bfd-template | | | | |
| Description | This command specifies the detect multiplier used for a BFD session. If a BFD control packet is not received for a period of <i>multiplier</i> x <i>receive-interval</i> , then the session is declared down. | | | | |
| Default | multiplier 3 | | | | |
| Parameters | <i>multiplier</i> — Specifies the multiplier. <table><tr><td>Values</td><td>3 to 20, integers</td></tr><tr><td>Default</td><td>3</td></tr></table> | Values | 3 to 20, integers | Default | 3 |
| Values | 3 to 20, integers | | | | |
| Default | 3 | | | | |

type

| | |
|--------------------|--|
| Syntax | [no] type cpm-np |
| Context | config>router>bfd>bfd-template |
| Description | This command selects the CPM network processor as the local termination point for the BFD session. This is enabled by default. |
| Default | no type |

single-sfm-overload

| | |
|--------------------|---|
| Syntax | single-sfm-overload [holdoff-time <i>holdoff-time</i>] no single-sfm-overload |
| Context | config>router |
| Description | This command configures OSPF, OSPFv3 and IS-IS to set overload when the router has fewer than the full set of SFMs functioning, which reduces forwarding capacity. Setting overload enables a router to still participate in exchanging routing information, but routes all traffic away from it. |

The conditions to set overload are as follows:

- 7750 SR-12/SR-7/SR-c12 and 7450 ESS-12/ESS-7/ESS-6 platforms: protocol sets overload if one of the SF/CPMs fails
- 7750 SR-12e and 7950 XRS platforms: protocol sets overload if two SFMs fail (two SFMs belonging to different SFM pairs on the XRS-40)

The **no** form of this command configures the router to not set overload if an SFM fails.

| | |
|-------------------|---|
| Default | no single-sfm-overload |
| Parameters | <i>holdoff-time</i> — This parameter specifies the delay between detecting SFM failures and setting overload. |
| Values | 1 to 600 seconds |
| Default | 0 seconds |

static-route-entry

| | | | | | | | | | | | | | | | | | | | |
|--------------------|---|-------------|-------------------------------|--------------------|---------|-------------|-------------------------------------|--|-------------------|--|----------------|--|---------------|--------------------|----------|-------------|-------------------------------|--------------------|---------|
| Syntax | static-route-entry { <i>ip-prefix/prefix-length</i> } [mcast] | | | | | | | | | | | | | | | | | | |
| Context | config>router | | | | | | | | | | | | | | | | | | |
| Description | <p>This command creates a static route entry for both the network and access routes. A prefix and netmask must be specified.</p> <p>Once the static route context for the specified prefix and netmask has been created, additional parameters associated with the static routes may be specified through the inclusion of additional static route parameter commands</p> <p>The no form of the command deletes the static route entry. If a static route needs to be removed when multiple static routes exist to the same destination, then as many parameters to uniquely identify the static route must be entered</p> <p>IPv6 static routes are not supported on the 7450 ESS except in mixed mode.</p> | | | | | | | | | | | | | | | | | | |
| Default | No static routes are defined. | | | | | | | | | | | | | | | | | | |
| Parameters | <p><i>ip-prefix/prefix-length</i> — The destination address of the static route.</p> <p>Values The following values apply to the 7750 SR and 7950 XRS:</p> <table><tr><td>ipv4-prefix</td><td>a.b.c.d (host bits must be 0)</td></tr><tr><td>ipv4-prefix-length</td><td>0 to 32</td></tr><tr><td>ipv6-prefix</td><td>x:x:x:x:x:x:x (eight 16-bit pieces)</td></tr><tr><td></td><td>x:x:x:x:x:d.d.d.d</td></tr><tr><td></td><td>x [0 to FFFF]H</td></tr><tr><td></td><td>d [0 to 255]D</td></tr><tr><td>ipv6-prefix-length</td><td>0 to 128</td></tr></table> <p>Values The following values apply to the 7450 ESS:</p> <table><tr><td>ipv4-prefix</td><td>a.b.c.d (host bits must be 0)</td></tr><tr><td>ipv4-prefix-length</td><td>0 to 32</td></tr></table> | ipv4-prefix | a.b.c.d (host bits must be 0) | ipv4-prefix-length | 0 to 32 | ipv6-prefix | x:x:x:x:x:x:x (eight 16-bit pieces) | | x:x:x:x:x:d.d.d.d | | x [0 to FFFF]H | | d [0 to 255]D | ipv6-prefix-length | 0 to 128 | ipv4-prefix | a.b.c.d (host bits must be 0) | ipv4-prefix-length | 0 to 32 |
| ipv4-prefix | a.b.c.d (host bits must be 0) | | | | | | | | | | | | | | | | | | |
| ipv4-prefix-length | 0 to 32 | | | | | | | | | | | | | | | | | | |
| ipv6-prefix | x:x:x:x:x:x:x (eight 16-bit pieces) | | | | | | | | | | | | | | | | | | |
| | x:x:x:x:x:d.d.d.d | | | | | | | | | | | | | | | | | | |
| | x [0 to FFFF]H | | | | | | | | | | | | | | | | | | |
| | d [0 to 255]D | | | | | | | | | | | | | | | | | | |
| ipv6-prefix-length | 0 to 128 | | | | | | | | | | | | | | | | | | |
| ipv4-prefix | a.b.c.d (host bits must be 0) | | | | | | | | | | | | | | | | | | |
| ipv4-prefix-length | 0 to 32 | | | | | | | | | | | | | | | | | | |

ip-address — The IP address of the IP interface. The *ip-addr* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

Values The following values apply to the 7750 SR and 7950 XRS:

| | |
|--------------|---|
| ipv4-address | a.b.c.d (host bits must be 0) |
| ipv6-address | x:x:x:x:x:x:x[-interface] x:x:x:x:x:x:d.d.d[-interface] x: [0..FFFF]H d: [0..255]D <i>interface</i> : 32 characters maximum, mandatory for link local addresses |

Values The following value applies to the 7450 ESS:

| | |
|--------------|-------------------------------|
| ipv4-address | a.b.c.d (host bits must be 0) |
|--------------|-------------------------------|

next-hop

| | | | | | | | |
|--------------------|--|-------------|-------------------|--------------|---------|--------------|---------------------------|
| Syntax | next-hop { <i>ip-address</i> <i>ip-int-name</i> <i>ipv6 address</i> } | | | | | | |
| Context | config>router>static-route-entry | | | | | | |
| Description | <p>This command specifies the directly connected next hop IP address or interface used to reach the destination. If the next hop is over an unnumbered interface or a point-to-point interface, the ip-int-name of the unnumbered or point-to-point interface (on this node) can be configured.</p> <p>If the next hop is over an unnumbered interface in the 7450 ESS router, the <i>ip-int-name</i> of the unnumbered interface (on this node) can be configured.</p> <p>The configured <i>ip-address</i> can be either on the network side or the access side on this node. The address must be associated with a network directly connected to a network configured on this node.</p> | | | | | | |
| Default | no next-hop | | | | | | |
| Parameters | <p><i>ip-address</i> <i>ip-int-name</i> <i>ipv6-address</i> — The IPv4/IPv6 address or interface of the next hop.</p> <p>Values The following values apply to the 7750 SR, 7450 ESS, and 7950 XRS:</p> <table><tbody><tr><td>ip-int-name</td><td>32 characters max</td></tr><tr><td>ipv4-address</td><td>a.b.c.d</td></tr><tr><td>ipv6-address</td><td>x:x:x:x:x:x-x[-interface]</td></tr></tbody></table> | ip-int-name | 32 characters max | ipv4-address | a.b.c.d | ipv6-address | x:x:x:x:x:x-x[-interface] |
| ip-int-name | 32 characters max | | | | | | |
| ipv4-address | a.b.c.d | | | | | | |
| ipv6-address | x:x:x:x:x:x-x[-interface] | | | | | | |

x:x:x:x:x:d.d.d.d[-interface]
x: [0..FFFF]H
d: [0..255]D
interface: 32 characters
maximum, mandatory for link
local addresses

IPv6 static routes are not supported on the 7450 ESS except in mixed mode.

indirect

| | | |
|-------------|--|---------------------------|
| Syntax | [no] indirect <i>ip-address</i> | |
| Context | config>router>static-route-entry | |
| Description | <p>This command specifies that the route is indirect and specifies the next hop IP address used to reach the destination.</p> <p>The configured <i>ip-address</i> is not directly connected to a network configured on this node. The destination can be reached via multiple paths. The indirect address can only resolved from dynamic routing protocol. Another static route cannot be used to resolve the indirect address.</p> <p>The <i>ip-address</i> configured here can be either on the network side or the access side and is typically at least one hop away from this node.</p> | |
| Default | no indirect | |
| Parameters | <i>ip-address</i> — The IP address of the IP interface. | |
| | Values | |
| | ipv4-address | a.b.c.d |
| | ipv6-address | x:x:x:x:x:x:x[-interface] |

black-hole

| | | |
|-------------|--|--|
| Syntax | [no] black-hole | |
| Context | config>router>static-route-entry | |
| Description | <p>This command specifies that the route is a black hole route. If the destination address on a packet matches this static route, it will be silently discarded.</p> | |
| Default | no black-hole | |

bfd-enable

| | |
|--------------------|---|
| Syntax | [no] bfd-enable |
| Context | config>router>static-route-entry>next-hop |
| Description | <p>This command associates the static route state to a BFD session between the local system and the configured nexthop.</p> <p>The remote end of the BFD session must also be configured to originate or accept the BFD session controlling the static route state.</p> <p>The no form of this command removes the association of the static route state to that of the BFD session.</p> |
| Default | no bfd-enable |

community

| | |
|--------------------|---|
| Syntax | [no] community <i>comm-id</i> |
| Context | config>router>static-route-entry>indirect config>router>static-route-entry>next-hop |
| Description | <p>This configuration option associates a BGP community with the static route. The community can be matched in route policies and is automatically added to BGP routes exported from the static route.</p> <p>The no form of this command removes the community association.</p> |
| Default | no community |
| Parameters | <i>comm-id</i> — Specifies community IDs, up to 72 characters. |
| Values | <i>[2 byte-as-number:comm-val well-known-comm]</i> where: <ul style="list-style-type: none">• <i>2 byte-as-number</i> — 0 to 65535• <i>comm-val</i> — 0 to 65535• <i>well-known-comm</i> — no-export no-export-subconfed no-advertise |

cpe-check

| | |
|----------------|--|
| Syntax | [no] cpe-check <i>cpe-ip-address</i> |
| Context | config>router>static-route-entry>next-hop config>router>static-route-entry>indirect |

| | |
|--------------------|---|
| Description | <p>This command enables CPE-check and specifies the IP address of the target CPE device.</p> <p>This option initiates a background ICMP ping test to the configured target IP address. The IP address can either be an IPv4 address for IPv4 static routes or an IPv6 address for IPv6 static routes. The target-ip-address cannot be in the same subnet as the static route subnet itself to avoid possible circular references. This option is mutually exclusive with BFD support on a given static route.</p> <p>The no form of this command disables the cpe-check option.</p> |
| Default | no cpe-check |
| Parameters | <i>cpe-ip-address</i> — Specifies the IP address of the CPE device. |

drop-count

| | |
|--------------------|---|
| Syntax | [no] drop-count <i>count</i> |
| Context | config>router>static-route-entry>next-hop>cpe-check config>router>static-route-entry>indirect>cpe-check |
| Description | This optional parameter specifies the number of consecutive ping-replies that must be missed to declare the CPE down and to deactivate the associated static route. |
| Default | drop-count 3 |
| Parameters | <i>count</i> — An integer count value. |
| Values | 1 to 255 |

interval

| | |
|--------------------|--|
| Syntax | [no] interval <i>seconds</i> |
| Context | config>router>static-route-entry>next-hop>cpe-check config>router>static-route-entry>indirect>cpe-check |
| Description | This optional parameter specifies the interval between ICMP pings to the target IP address. |
| Default | interval 1 |
| Parameters | <i>seconds</i> — An integer interval value. |
| Values | 1 to 255 |

padding-size

| | |
|---------------|--|
| Syntax | [no] padding-size <i>padding-size</i> |
|---------------|--|

| | |
|--------------------|---|
| Context | config>router>static-route-entry>next-hop>cpe-check config>router>static-route-entry>indirect>cpe-check |
| Description | This optional parameter specifies the amount of padding to add to the ICMP packet in bytes. The parameter is only applicable when the cpe-check option is used with the associated static route. |
| Default | padding-size 56 |
| Parameters | <i>padding-size</i> — An integer value. |
| Values | 0 to 16384 bytes |

log

| | |
|--------------------|--|
| Syntax | [no] log |
| Context | config>router>static-route-entry>next-hop>cpe-check config>router>static-route-entry>indirect>cpe-check |
| Description | This optional parameter enables the ability to log transitions between active and in-active based on the CPE connectivity check. Events will be sent to the system log, syslog and SNMP traps. |
| Default | no log |

description

| | |
|--------------------|--|
| Syntax | [no] description <i>description-string</i> |
| Context | config>router>static-route-entry>next-hop config>router>static-route-entry>indirect config>router>static-route-entry>black-hole |
| Description | This command creates a text description stored in the configuration file for a configuration context. The no form of the command removes the description string from the context |
| Default | no description |
| Parameters | <i>description-string</i> — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. |

destination-class

| | |
|--------------------|---|
| Syntax | [no] destination-class <i>dest-index</i> |
| Context | config>router>static-route-entry>next-hop config>router>static-route-entry>indirect config>router>static-route-entry>black-hole |
| Description | This command configures the policy accounting destination-class index to be used when incrementing accounting statistic for traffic matching the associated static route. The no form of the command removes the associated destination-class from the associated static route nexthop. |
| Default | no destination-class |
| Parameters | <i>dest-index</i> — The destination index integer value. Values 1 to 255 |

dynamic-bgp

| | |
|--------------------|--|
| Syntax | [no] dynamic-bgp |
| Context | config>router>static-route-entry>black-hole |
| Description | This optional command controls the behavior of the associated static route so that if a matching BGP route to the same exact prefix is present in BGP, the static route's nexthop is set to the BGP's nexthop value. If there is no matching active BGP route, the static route's nexthop is set to be a black-hole nexthop. |
| Default | no dynamic-bgp |

generate-icmp

| | |
|--------------------|---|
| Syntax | [no] generate-icmp |
| Context | config>router>static-route-entry>black-hole |
| Description | This optional command causes the ICMP unreachable messages to be sent when received packets match the associated static route. By default, the ICMP unreachable messages for those types of static routes are not generated. This command can only be associated with a static route that has a blackhole next-hop The no form of this command removes the black-hole nexthop from the static route configuration. |
| Default | no generate-icmp |

forwarding-class

| | |
|--------------------|---|
| Syntax | [no] forwarding-class {be l2 af l1 h2 ef h1 nc} |
| Context | config>router>static-route-entry>next-hop config>router>static-route-entry>indirect config>router>static-route-entry>next-hop |
| Description | This command specifies the enqueueing forwarding class that should be associated with traffic matching the associate static route. If this parameter is not specified, the packet will use the forwarding-class association based on default classification or other QoS Policy associations. |
| Default | no forwarding-class |
| Parameters | be l2 af l1 h2 ef h1 nc — Specifies the forwarding class. Values be l2 af l1 h2 ef h1 nc |

ldp-sync

| | |
|--------------------|--|
| Syntax | [no] ldp-sync |
| Context | config>router>static-route-entry>next-hop config>router>static-route-entry>indirect config>router>static-route-entry>next-hop |
| Description | <p>This command extends the LDP synchronization feature to a static route. When an interface comes back up, it is possible that a preferred static route using the interface as next-hop for a given prefix is enabled before the LDP adjacency to the peer LSR comes up on this interface. In this case, traffic on an SDP that uses the static route for the far-end address would be black-holed until the LDP session comes up and the FECs exchanged.</p> <p>This option when enabled delays the activation of the static route until the LDP session comes up over the interface and the ldp-sync-timer configured on that interface has expired</p> |
| Default | no ldp-sync |

metric

| | |
|----------------|--|
| Syntax | [no] metric <i>metric-value</i> |
| Context | config>router>static-route-entry>next-hop config>router>static-route-entry>indirect |

Description This command specifies the cost metric for the static route, expressed as a decimal integer. This value is used when importing the static route into other protocols such as OSPF. When the metric is configured as 0 then the metric configured in OSPF, default-import-metric, applies. When modifying the metric of an existing static route, the preference will not change unless specified. This value is also used to determine which static route to install in the forwarding table.

If there are multiple static routes with the same preference but different metrics then the lower cost (metric) route will be installed.

The **no** form of this command returns the metric to the default value

Default metric 1

Parameters *metric-value* — Specifies the cost metric value.

Values 0 to 65535

preference

Syntax **[no] preference** *preference-value*

Context config>router>static-route-entry>next-hop
config>router>static-route-entry>indirect
config>router>static-route-entry>black-hole

Description This command specifies the route preference to be assigned to the associated static route. The lower the preference value the more preferred the route is considered.

[Table 7](#) shows the default route preference based on the route source.

Table 7 Default Route Preference

| Label | Preference | Configurable |
|------------------------|------------|--------------|
| Direct attached | 0 | No |
| Static route | 5 | Yes |
| OSPF Internal routes | 10 | Yes |
| IS-IS level 1 internal | 15 | Yes |
| IS-IS level 2 internal | 18 | Yes |
| OSPF external | 150 | Yes |
| IS-IS level 1 external | 160 | Yes |
| IS-IS level 2 external | 165 | Yes |
| BGP | 170 | Yes |

The **no** form of this command returns the associated static route preference to its default value.

| | |
|-------------------|---|
| Default | preference 5 |
| Parameters | <i>preference-value</i> — Specifies the route preference value. |
| Values | 1 to 255 |

prefix-list

| | |
|--------------------|--|
| Syntax | [no] prefix-list <i>name</i> { all none any } |
| Context | config>router>static-route-entry>next-hop config>router>static-route-entry>indirect config>router>static-route-entry>black-hole |
| Description | This command associates a new constraint to the associated static route such that the static route is only active if any , none , or all of the routes in the prefix list are present and active in the route-table. |
| Default | no prefix-list |
| Parameters | <i>name</i> — Specifies the name of a currently configured prefix-list. all — Specifies that the static route condition is met if all prefixes in the prefix-list must be present in the active route-table. none — Specifies that the static route condition is met if none of the prefixes in the named prefix-list can be present in the active route-table. any — Specifies that the static route condition is met if any prefixes in the prefix-list are present in the active route-table. |

priority

| | |
|--------------------|--|
| Syntax | [no] priority { low high } |
| Context | config>router>static-route-entry>next-hop>forwarding-class config>router>static-route-entry>indirect>forwarding-class |
| Description | This optional command associates an enqueueing priority with the static route. The options are either high or low, with low being the default. This parameter has the ability to affect the likelihood that a packet will be enqueued at SAP ingress in the face of ingress congestion. Once a packet is enqueued into an ingress buffer, the significance of this parameter is lost. |
| Default | priority low |

- Parameters**
- low** — Setting the enqueueing parameter for a packet to **low** decreases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing. Once the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost.
- high** — Setting the enqueueing parameter for a packet to **high** increases the likelihood of enqueueing the packet when the ingress queue is congested. Ingress enqueueing priority only affects ingress SAP queuing. Once the packet is placed in a buffer on the ingress queue, the significance of the enqueueing priority is lost..

shutdown

- Syntax** **[no] shutdown**
- Context** config>router>static-route-entry>black-hole
config>router>static-route-entry>indirect
config>router>static-route-entry>next-hop
- Description** This command causes the static route to be placed in an administratively down state and removed from the active route-table
- Default** no shutdown

source-class

- Syntax** **[no] source-class** *source-index*
- Context** config>router>static-route-entry>indirect
config>router>static-route-entry>next-hop
- Description** This command configures the policy accounting source-class index to be used when incrementing accounting statistic for traffic matching the associated static route.
- If source route policy accounting is enabled and a source-class index is configured, traffic with a source IP address matches the associated static route, the source accounting statistics for the specified class will be incremented.
- The **no** form of the command removes the associated destination-class from the associated static route nexthop.
- Default** no source-class
- Parameters** *source-index* — Specifies an integer value for the accounting source class index.
- Values** 1 to 255

tag

| | |
|--------------------|---|
| Syntax | [no] tag <i>tag-value</i> |
| Context | config>router>static-route-entry>indirect config>router>static-route-entry>next-hop |
| Description | <p>This command adds a 32-bit integer tag to the associated static route.</p> <p>The tag value can be used in route policies to control distribution of the route into other protocols.</p> |
| Default | no tag |
| Parameters | <i>tag-value</i> — Specifies an integer tag value. |
| Values | 32 bit integer |

tunnel-next-hop

| | |
|--------------------|---|
| Syntax | tunnel-next-hop |
| Context | config>router>static-route-entry>indirect |
| Description | This command enables the context to configure the static route's nexthop to be resolved to an indirect tunnel next-hop. |

disallow-igp

| | |
|--------------------|---|
| Syntax | [no] disallow-igp |
| Context | config>router>static-route-entry>indirect>tunnel-next-hop |
| Description | <p>This optional command determines if the associated static route can be resolved via an IGP next-hop in the RTM if no tunnel next-hops are found in TTM.</p> <p>When configured, the associated static route will not be resolved to an available IGP route in the RTM.</p> <p>The no form of the command returns the behavior to the default, which does allow for the static route to be resolved via an IGP route in the RTM if no tunnel next-hop can be found in the TTM.</p> |
| Default | no disallow-igp |

resolution

| | |
|--------------------|--|
| Syntax | resolution {any disabled filter} no resolution |
| Context | config>router>static-route-entry>indirect>tunnel-next-hop |
| Description | This command determines how the associated static route can be resolved to a tunnel next-hop. |
| Default | resolution any |
| Parameters | <p>any — Allows the associated static route to be resolved to any active entry in the TTM, following the TTM preference order.</p> <p>disabled — Disables the associated static route to be resolved to any active entry in the TTM. As a result, the static route can only be resolved via IP RTM resolution of the static route's nexthop.</p> <p>filter — Allows the associated static route to be resolved to active tunnels in the TTM using the resolution-filter restrictions.</p> |

resolution-filter

| | |
|--------------------|---|
| Syntax | resolution-filter |
| Context | config>router>static-route-entry>indirect>tunnel-next-hop |
| Description | <p>This command creates the context to specify the tunnel next-hop resolution options.</p> <p>If one or more tunnel filter criteria are specified, the static route nexthop will be resolved to an available tunnel from one of those LSP types. The tunnel type will be selected following the TTM preference.</p> |

ldp

| | |
|--------------------|--|
| Syntax | [no] ldp |
| Context | config>router>static-route-entry>indirect>tunnel-next-hop>resolution-filter |
| Description | This command enables the use of LDP sourced tunnel entries in the TTM to resolve the associated static route next-hop. |
| Default | no ldp |

rsvp-te

| | |
|--------------------|---|
| Syntax | [no] rsvp-te |
| Context | config>router>static-route-entry>indirect>tunnel-next-hop>resolution-filter |
| Description | <p>This command enables the use of RSVP-TE sourced tunnel entries in the TTM to resolve the associated static route next-hop.</p> <p>The rsvp-te value instructs the code to search for the set of lowest metric RSVP-TE LSPs to the address of the indirect next-hop. The LSP metric is provided by MPLS in the tunnel table. The static route treats a set of RSVP-TE LSPs with the same lowest metric as an ECMP set. The user has the option of configuring a list of RSVP-TE LSP names to be used exclusively instead of searching in the tunnel table. In that case, all LSPs must have the same LSP metric in order for the static route to use them as an ECMP set. Otherwise, only the LSPs with the lowest common metric value will be selected.</p> <p>A P2P auto-lsp that is instantiated via an LSP template can be selected in TTM when resolution is set to any. However, it is not recommended to configure an auto-lsp name explicitly under the rsvp-te node as the auto-generated name can change if the node reboots, which will blackhole the traffic of the static route.</p> |
| Default | no rsvp-te |

lsp

| | |
|--------------------|---|
| Syntax | [no] lsp <i>name</i> |
| Context | config>router>static-route-entry>indirect>tunnel-next-hop>resolution-filter>rsvp-te config>router>static-route-entry>indirect>tunnel-next-hop>resolution-filter>sr-te |
| Description | This command restricts the search for a resolving LSP to a specific set of named LSPs. Only those LSPs named in the associated name list will be searched for a match to resolve the associated static route. |
| Parameters | <i>name</i> — Specifies the name of the LSP to be searched for a valid resolving tunnel for the static route's next-hop. |

sr-ospf

| | |
|--------------------|--|
| Syntax | [no] sr-ospf |
| Context | config>router>static-route-entry>indirect>tunnel-next-hop>resolution-filter |
| Description | This command enables the use of sr-ospf sourced tunnel entries in the TTM to resolve the associated static route next-hop. |
| Default | no sr-ospf |

sr-isis

| | |
|--------------------|--|
| Syntax | [no] sr-isis |
| Context | config>router>static-route-entry>indirect>tunnel-next-hop>resolution-filter |
| Description | This command enables the use of sr-isis sourced tunnel entries in the TTM to resolve the associated static route next-hop. |
| Default | no sr-isis |

sr-te

| | |
|--------------------|---|
| Syntax | [no] sr-te |
| Context | config>router>static-route-entry>indirect>tunnel-next-hop>resolution-filter |
| Description | The sr-te value instructs the code to search for the set of lowest metric SR-TE LSPs to the address of the indirect next-hop. The LSP metric is provided by MPLS in the tunnel table. The static route treats a set of SR-TE LSPs with the same lowest metric as an ECMP set. The user has the option of configuring a list of SR-TE LSP names to be used exclusively instead of searching in the tunnel table. In that case, all LSPs must have the same LSP metric in order for the static route to use them as an ECMP set. Otherwise, only the LSPs with the lowest common metric value are selected. |
| Default | no sr-te |

validate-next-hop

| | |
|--------------------|---|
| Syntax | [no] validate-next-hop |
| Context | config>router>static-route-entry>next-hop |
| Description | <p>This optional command tracks the state of the next-hop in the IPv4 ARP cache or IPv6 Neighbor Cache. When the next-hop is not reachable and is removed from the ARP or Neighbor Cache, the next-hop will no longer be considered valid and the associated static-route state removed from the active route-table.</p> <p>When the next-hop is reachable again and present in the ARP/Neighbor Cache, the static route will be considered valid and is subject to being placed into the active route-table.</p> |
| Default | no validate-next-hop |

disallow-igp

| | |
|---------------|---------------------|
| Syntax | disallow-igp |
|---------------|---------------------|

no disallow-igp

| | |
|--------------------|---|
| Context | config>router>static-route-entry>tunnel-next-hop |
| Description | This command is for indirect static routes using tunnel next-hops. When enabled, the static route will not be activated using IGP next-hops in RTM if no tunnel next-hops are found in TTM. |
| Default | no disallow-igp |

triggered-policy

| | |
|--------------------|---|
| Syntax | triggered-policy no triggered-policy |
| Context | config>router |
| Description | <p>This command triggers route policy re-evaluation.</p> <p>By default, when a change is made to a policy in the config router policy options context and then committed, the change is effective immediately. There may be circumstances when the changes should or must be delayed; for example, if a policy change is implemented that would affect every BGP peer on a router, the consequences could be dramatic. It would be more effective to control changes on a peer-by-peer basis.</p> <p>If the triggered-policy command is enabled, and a given peer is established, and you want the peer to remain up, in order for a change to a route policy to take effect, a clear command with the <i>soft</i> or <i>soft inbound</i> option must be used; for example, clear router bgp neighbor x.x.x.x soft. This keeps the peer up, and the change made to a route policy is applied only to that peer or group of peers.</p> |
| Default | no triggered-policy |

ttl-propagate

| | |
|--------------------|---|
| Syntax | ttl-propagate |
| Context | config>router |
| Description | This command enables the context to configure TTL propagation for transit and locally generated packets in the Global Routing Table (GRT) and VPRN routing contexts |

label-route-local

| | |
|----------------|---------------------------------------|
| Syntax | label-route-local [all none] |
| Context | config>router>ttl-propagate |

| | |
|--------------------|--|
| Description | <p>This command configures the TTL propagation for locally generated packets which are forwarded over a BGP label route in the Global Routing Table (GRT) context.</p> <p>For IPv4 and IPv6 packets forwarded using a RFC 3107 label route in the global routing instance, including 6PE, the all value of the command enables TTL propagation from the IP header into all labels in the transport label stack. The none value reverts to the default mode which disables TTL propagation from the IP header to the labels in the transport label stack. This command does not have a no version.</p> <p>The TTL of the IP packet is always propagated into the RFC 3107 label itself, and this command only controls the propagation into the transport labels, for example, labels of the RSVP or LDP LSP to which the BGP label route resolves and which are pushed on top of the BGP label.</p> <p>If the BGP peer advertised the implicit-null label value for the BGP label route, the TTL propagation will not follow the configuration described, but will follow the configuration to which the BGP label route resolves:</p> <p>RSVP LSP shortcut:</p> <ul style="list-style-type: none"> • configure router mpls shortcut-local-ttl-propagate <p>LDP LSP shortcut:</p> <ul style="list-style-type: none"> • configure router ldp shortcut-local-ttl-propagate <p>This feature does not impact packets forwarded over BGP shortcuts. The ingress LER operates in uniform mode by default and can be changed into pipe mode using the configuration of TTL propagation for RSVP or LDP LSP shortcut listed.</p> |
| Default | label-route-local none |
| Parameters | <p>none — The TTL of the IP packet is not propagated into the transport label stack.</p> <p>all — The TTL of the IP packet is propagated into all labels of the transport label stack.</p> |

label-route-transit

| | |
|--------------------|---|
| Syntax | label-route-transit [all none] |
| Context | config>router>ttn-propagate |
| Description | <p>This command configures the TTL propagation for transit packets which are forwarded over a BGP label route in the Global Routing Table (GRT) context.</p> <p>For IPv4 and IPv6 packets forwarded using a RFC 3107 label route in the global routing instance, including 6PE, the all value of the command enables TTL propagation from the IP header into all labels in the transport label stack. The none value reverts to the default mode which disables TTL propagation from the IP header to the labels in the transport label stack. This command does not have a no version.</p> |

The TTL of the IP packet is always propagated into the RFC 3107 label itself, and this command only controls the propagation into the transport labels, for example, labels of the RSVP or LDP LSP to which the BGP label route resolves and which are pushed on top of the BGP label.

If the BGP peer advertised the implicit-null label value for the BGP label route, the TTL propagation will not follow the configuration described, but will follow the configuration to which the BGP label route resolves.

RSVP LSP shortcut:

- configure router mpls shortcut-transit-ttl-propagate

LDP LSP shortcut:

- configure router ldp shortcut-transit-ttl-propagate

This feature does not impact packets forwarded over BGP shortcuts. The ingress LER operates in uniform mode by default and can be changed into pipe mode using the configuration of TTL propagation for the listed RSVP or LDP LSP shortcut.

| | |
|-------------------|---|
| Default | label-route-transit none |
| Parameters | none — The TTL of the IP packet is not propagated into the transport label stack. all — The TTL of the IP packet is propagated into all labels of the transport label stack. |

lsr-label-route

| | |
|--------------------|---|
| Syntax | lsr-label-route [all none] |
| Context | config>router>ttl-propagate |
| Description | <p>This command configures the TTL propagation for transit packets at a router acting as an LSR for a BGP label route.</p> <p>When an LSR swaps the BGP label for a ipv4 prefix packet, therefore acting as a ABR, ASBR, or data-path Route-Reflector (RR) in the base routing instance, or swaps the BGP label for a vpn-ipv4 or vpn-ipv6 prefix packet, therefore acting as an inter-AS Option B VPRN ASBR or VPRN data path Route-Reflector (RR), the all value of this command enables TTL propagation of the decremented TTL of the swapped BGP label into all outgoing LDP or RSVP transport labels.</p> <p>When an LSR swaps a label or stitches a label, it always writes the decremented TTL value into the outgoing swapped or stitched label. What this feature controls is whether this decremented TTL value is also propagated to the transport label stack pushed on top of the swapped or stitched label.</p> |

The none value reverts to the default mode which disables TTL propagation. This changes the existing default behavior which propagates the TTL to the transport label stack. When a customer upgrades, the new default becomes in effect. This command does not have a no version.

This feature also controls the TTL propagation at an LDP-BGP stitching LSR in the LDP to BGP stitching direction. It also controls the TTL propagation in Carrier Supporting Carrier (CsC) VPRN at both the CsC CE and CsC PE.

SR OS does not support ASBR or data path RR functionality for labeled IPv6 routes in the global routing instance (6PE). As such the CLI command of this feature has no impact on prefix packets forwarded in this context.

| | |
|-------------------|---|
| Default | lsr-label-route none |
| Parameters | none — The TTL of the swapped label is not propagated into the transport label stack. all — The TTL of the swapped label is propagated into all labels of the transport label stack. |

vprn-local

| | |
|--------------------|--|
| Syntax | vprn-local [all vc-only none] |
| Context | config>router>ttn-propagate |
| Description | <p>This command configures the TTL propagation for locally generated packets which are forwarded over a MPLS LSPs in all VPRN service contexts.</p> <p>For vpn-ipv4 and vpn-ipv6 packets forwarded in the context of all VPRN services in the system, including 6VPE packets, the all value of the command enables TTL propagation from the IP header into all labels in the stack:</p> <p>The user can enable the TTL propagation behavior separately for locally generated packets by CPM (vprn-local) and for user and control packets in transit at the node (vprn-transit).</p> <p>The vc-only value reverts to the default behavior by which the IP TTL is propagated into the VC label but not to the transport labels in the stack. The user can explicitly set the default behavior by configuring the vc-only value. This command does not have a no version.</p> <p>The value none allows the user to disable the propagation of the IP TTL to all labels in the stack, including the VC label. This is needed for a transparent operation of UDP traceroute in VPRN inter-AS option B such that the ingress and egress ASBR nodes are not traced.</p> <p>The user can override the global configuration within each VPRN instance using the following commands:</p> <ul style="list-style-type: none"> • config service vprn ttl-propagate local [inherit none vc-only all] • config service vprn ttl-propagate transit [inherit none vc-only all] |

The default behavior for a given VPRN instance is to inherit the global configuration for the same command. The user can explicitly set the default behavior by configuring the inherit value.

When a packet is received in a VPRN context but is looked up in the Global Routing Table (GRT), for example, leaking to GRT is enabled, the behavior of the TTL propagation is governed by the RSVP or LDP shortcut configuration when the matching routing is a LSP shortcut route. It is governed by the BGP label route configuration when the matching route is a RFC 3107 label route or a 6PE route.

When a packet is received on one VPRN instance and is redirected using Policy Based Routing (PBR) to be forwarded in another VPRN instance, the TTL propagation is governed by the configuration of the outgoing VPRN instance.

| | |
|-------------------|---|
| Default | vprn-local vc-only |
| Parameters | <p>none — The TTL of the IP packet is not propagated into the VC label or labels in the transport label stack</p> <p>vc-only — The TTL of the IP packet is propagated into the VC label and not into the labels in the transport label stack.</p> <p>all — The TTL of the IP packet is propagated into the VC label and all labels in the transport label stack.</p> |

vprn-transit

| | |
|--------------------|--|
| Syntax | vprn-transit [all vc-only none] |
| Context | config>router>ttd-propagate |
| Description | <p>This command configures the TTL propagation for in transit packets which are forwarded over a MPLS LSPs in all VPRN service contexts. For vpn-ipv4 and vpn-ipv6 packets forwarded in the context of all VPRN services in the system, including 6VPE packets, the all value of the command enables TTL propagation from the IP header into all labels in the stack:</p> <p>The user can enable the TTL propagation behavior separately for locally generated packets by CPM (vprn-local) and for user and control packets in transit at the node (vprn-transit).</p> <p>The vc-only value reverts to the default behavior by which the IP TTL is propagated into the VC label but not to the transport labels in the stack. The user can explicitly set the default behavior by configuring the vc-only value. This command does not have a no version.</p> <p>The value none allows the user to disable the propagation of the IP TTL to all labels in the stack, including the VC label. This is needed for a transparent operation of UDP trace-route in VPRN inter-AS option B such that the ingress and egress ASBR nodes are not traced.</p> <p>The user can override the global configuration within each VPRN service instance using the following commands:</p> |

- config service vprn ttl-propagate local [inherit | none | vc-only | all]

- config service vprn ttl-propagate transit [inherit | none | vc-only | all]

The default behavior for a given VPRN instance is to inherit the global configuration for the same command. The user can explicitly set the default behavior by configuring the inherit value.

When a packet is received in a VPRN context but is looked up in the Global Routing Table (GRT), for example, leaking to GRT is enabled, the behavior of the TTL propagation is governed by the RSVP or LDP shortcut configuration when the matching routing is a LSP shortcut route. It is governed by the BGP label route configuration when the matching route is a RFC 3107 label route or a 6PE route.

When a packet is received on one VPRN instance and is redirected using Policy Based Routing (PBR) to be forwarded in another VPRN instance, the TTL propagation is governed by the configuration of the outgoing VPRN instance

| | |
|-------------------|---|
| Default | vprn-transit vc-only |
| Parameters | <p>none — The TTL of the IP packet is not propagated into the VC label or labels in the transport label stack</p> <p>vc-only — The TTL of the IP packet is propagated into the VC label and not into the labels in the transport label stack.</p> <p>all — The TTL of the IP packet is propagated into the VC label and all labels in the transport label stack.</p> |

2.12.2.3 Router L2TP Commands

Router L2TP commands only apply to the 7750 SR and 7450 ESS.

l2tp

| | |
|--------------------|---|
| Syntax | l2tp |
| Context | config>router |
| Description | This command enables the context to configure L2TP parameters. L2TP extends the PPP model by allowing Layer 2 and PPP endpoints to reside on different devices interconnected by a packet-switched network. |

calling-number-format

| | |
|---------------|---|
| Syntax | calling-number-format <i>ascii-spec</i> no calling-number-format |
|---------------|---|

| | | | |
|--------------------|---|--------------------|--|
| Context | config>router>l2tp | | |
| Description | This command what string to put in the Calling Number AVP, for L2TP control messages related to a session in this L2TP protocol instance. | | |
| Default | calling-number-format "%S %s" | | |
| Parameters | <i>ascii-spec</i> — Specifies the L2TP calling number AVP. | | |
| | Values | | |
| | ascii-spec | char-specification | ascii-spec |
| | | char-specification | ascii-char char-origin |
| | | ascii-char | a printable ASCII character |
| | | char-origin | %origin |
| | | origin | S c r s l |
| | | | S system name, the value of TIMETRA-CHASSIS- MIB::tmnxChassisName |
| | | | c Agent Circuit Id |
| | | | r Agent Remote Id |
| | | | s SAP ID, formatted as a character string |
| | | | l Logical Line ID |

eth-tunnel

| | |
|--------------------|--|
| Syntax | eth-tunnel |
| Context | config>router>l2tp |
| Description | This command enables the context to configure Ethernet tunnel client parameters. |

reconnect-timeout

| | |
|--------------------|--|
| Syntax | reconnect-timeout <i>seconds</i> no reconnect-timeout |
| Context | config>router>l2tp>eth-tunnel |
| Description | This command configures the number of seconds that the Ethernet tunnel client of L2TPv3 waits before attempting to re-establish a new session after a session setup fails or a session closes. |

The **no** form of the command returns **reconnect-timeout** to an infinite timeout value, meaning that reconnection will not be attempted by the local client.

| | |
|-------------------|--|
| Default | no reconnect-timeout (infinite timeout) |
| Parameters | <i>seconds</i> — Specifies the number of seconds before a session reconnection is attempted after a previous session or session setup fails. |
| Values | 10 to 3600 |

exclude-avps

| | |
|--------------------|---|
| Syntax | exclude-avps <i>calling-number</i> no exclude-avps |
| Context | config>router>l2tp |
| Description | This command configures the L2TP AVPs to exclude. |
| Default | no exclude-avps |

group-session-limit

| | |
|--------------------|--|
| Syntax | group-session-limit <i>session-limit</i> group-session-limit unlimited no group-session-limit |
| Context | config>router>l2tp |
| Description | This command configures the L2TP session limit for each group of this router. |
| Default | no group-session-limit |

l2tpv3

| | |
|--------------------|--|
| Syntax | l2tpv3 |
| Context | config>router>l2tp config>router>l2tp>group |
| Description | This command enables the context to configure L2TPv3 parameters. |

cookie-length

| | |
|---------------|--|
| Syntax | cookie-length {4 8 default} |
|---------------|--|

no cookie-length

| | |
|--------------------|---|
| Context | config>router>l2tp>l2tpv3 config>router>l2tp>group>l2tpv3 |
| Description | This command configures the length of the optional cookie field. The default parameter only applies in the config>router>l2tp>group>l2tpv3 context. The no form of the command returns the cookie-length to a default of none . |
| Default | no cookie-length |
| Parameters | 4 — Specifies the cookie length as 4 bytes. 8 — Specifies the cookie length as 8 bytes. default — In the config>router>l2tp>group>l2tpv3 , this parameter references the cookie length configured in the config>router>l2tp>l2tpv3 context. |

digest-type

| | |
|--------------------|---|
| Syntax | digest-type {default none md5 sha1} no digest-type |
| Context | config>router>l2tp>l2tpv3 config>router>l2tp>group>l2tpv3 |
| Description | This command configures the hashing algorithm used to calculate the message digest. The default parameter only applies in the config>router>l2tp>group>l2tpv3 context. The no form of the command returns the digest-type to none . |
| Default | no digest-type |
| Parameters | default — In the config>router>l2tp>group>l2tpv3 , this parameter references the digest type configured in the config>router>l2tp>l2tpv3 context. none — Specifies that no digest should be used. md5 — Specifies that the MD5 algorithm should be used. sha1 — Specifies that the SHA1 algorithm should be used. |

nonce-length

| | |
|----------------|--|
| Syntax | nonce-length {length default} no nonce-length |
| Context | config>router>l2tp>l2tpv3 config>router>l2tp>group>l2tpv3 |

| | |
|--------------------|---|
| Description | <p>This command configures the length for the local L2TPv3 nonce (random number) value used in the Nonce AVP. The default parameter only applies in the config>router>l2tp>group>l2tpv3 context</p> <p>The no form of the command returns the nonce-length to a default of none.</p> |
| Default | no nonce-length |
| Parameters | <p><i>length</i> — Specifies the length of the Nonce AVP value.</p> <p>Values 16 to 64</p> <p>default — In the config>router>l2tp>group>l2tpv3, this parameter references the nonce length configured in the config>router>l2tp>l2tpv3 context.</p> |

private-tcp-mss-adjust

| | |
|--------------------|--|
| Syntax | <p>private-tcp-mss-adjust <i>octets</i></p> <p>no private-tcp-mss-adjust</p> |
| Context | <p>config>router>l2tp>l2tpv3</p> <p>config>service>vprn>l2tp>l2tpv3</p> |
| Description | <p>This command enables TCP MSS adjust for L2TPv3 tunnels on the private side of the service level. When this command is configured, the system updates the TCP MSS option value of the received TCP SYN packet on the private side.</p> <p>Note that this command can be overridden by the corresponding configuration on the group or tunnel level.</p> <p>The no form of this command disables TCP MSS adjust on the private side.</p> |
| Default | no private-tcp-mcc-adjust |
| Parameters | <p><i>octets</i> — Specifies the new TCP MSS value in octets.</p> <p>Values 512 to 9000</p> |

public-tcp-mss-adjust

| | |
|--------------------|---|
| Syntax | <p>public-tcp-mss-adjust <i>octets</i></p> <p>no public-tcp-mss-adjust</p> |
| Context | <p>config>router>l2tp>l2tpv3</p> <p>config>service>vprn>l2tp>l2tpv3</p> |
| Description | <p>This command enables TCP MSS adjust for L2TPv3 tunnels on the public side on the service level. When the command is configured, the system updates the TCP MSS option value of the received TCP SYN packet on the public side that is encapsulated in the L2TPv3 tunnel.</p> |

Note that this command can be overridden by the corresponding configuration on the group or tunnel level.

The **no** form of this command disables TCP MSS adjust on the public side.

| | |
|-------------------|---|
| Default | no public-tcp-mss-adjust |
| Parameters | <i>octets</i> — Specifies the new TCP MSS value in octets |
| Values | 512 to 9000 |

private-tcp-mss-adjust

| | |
|--------------------|---|
| Syntax | private-tcp-mss-adjust <i>octets</i> private-tcp-mss-adjust default no private-tcp-mss-adjust |
| Context | config>router>l2tp>group>l2tpv3 config>service>vpn>l2tp>group>l2tpv3 |
| Description | <p>This command enables TCP MSS adjust for L2TPv3 tunnels on the private side of the group or tunnel level. When this command is configured, the system updates the TCP MSS option value of the received TCP SYN packet on the private side.</p> <p>With the default parameter, the system uses the upper-level configuration. With the non-default parameter, the system uses this configuration instead of the upper level configuration.</p> <p>The no form of this command disables TCP MSS adjust on the private side.</p> |
| Default | no private-tcp-mcc-adjust |
| Parameters | <i>octets</i> — Specifies the new TCP MSS value in octets. Values 512 to 9000 default — Specifies to use the upper-level configuration |

public-tcp-mss-adjust

| | |
|----------------|---|
| Syntax | public-tcp-mss-adjust <i>octets</i> public-tcp-mss-adjust default no public-tcp-mss-adjust |
| Context | config>router>l2tp>group>l2tpv3 config>service>vpn>l2tp>group>l2tpv3 |

| | |
|--------------------|---|
| Description | <p>This command enables TCP MSS adjust for L2TPv3 tunnels on the public side on the group or tunnel level. When the command is configured, the system updates the TCP MSS option value of the received TCP SYN packet on the public side that is encapsulated in the L2TPv3 tunnel.</p> <p>With the default parameter, the system uses the upper level configuration. With the non-default parameter, the system uses this configuration instead of the upper level configuration.</p> <p>The no form of this command disables TCP MSS adjust on the public side.</p> |
| Default | no public-tcp-mss-adjust |
| Parameters | <i>octets</i> — Specifies the new TCP MSS value in octets |
| Values | 512 to 9000 |
| | default — Specifies to use the upper-level configuration |

rem-router-id

| | |
|--------------------|---|
| Syntax | rem-router-id <i>ip-addr</i> no rem-router-id |
| Context | config>router>l2tp>group>l2tpv3 |
| Description | <p>This command configures the IP address that should be used within the Remote Router-ID AVP.</p> <p>The no form of this command removes the configured IP address.</p> |
| Default | no rem-router-id |
| Parameters | <i>ip-addr</i> — Specifies an IP address to be used within the Remote Router-ID AVP. |

pw-cap-list

| | |
|--------------------|---|
| Syntax | pw-cap-list { ethernet ethernet-vlan } no pw-cap-list |
| Context | config>router>l2tp>group>l2tpv3 |
| Description | <p>This command configures the allowable pseudowire capability list that is advertised to the far end. An empty list results in both pseudowire capabilities being advertised.</p> <p>The no form of this command removes the list and advertises both pseudowire capabilities to the far end.</p> |
| Default | no pw-cap-list |

-
- Parameters** **ethernet** — Specifies that the Ethernet pseudo-wire type is advertised.
- ethernet-vlan** — Specifies that the Ethernet-VLAN pseudo-wire type is advertised. This parameter is only supported in SR OS Release 14.0 R4 or later.

track-password-change

- Syntax** **[no] track-password-change**
- Context** config>router>l2tp>group>l2tpv3
- Description** This command enables tracking of password changes, allowing password tunnel passwords to be changed without bringing down active tunnels or sessions. This is only supported with L2TPv3.
- The **no** form of the command disables password change tracking.
- Default** no track-password-change

transport-type

- Syntax** **transport-type ip**
 no transport-type
- Context** config>router>l2tp>l2tpv3
- Description** This command configures the transport type to be used to carry the L2TPv3 tunnel. Currently, only IP transport is supported.
- The **no** form of this command returns the **transport-type** to the default value.
- Default** no transport-type
- Parameters** **ip** — Specifies that IP should be used as the transport type for the L2TPv3 tunnel.

next-attempt

- Syntax** **next-attempt {same-preference-level | next-preference-level}**
 no next-attempt
- Context** config>router>l2tp
 config>service>vpn>l2tp
- Description** This command enables tunnel selection algorithm based on the tunnel preference level.

| | |
|-------------------|--|
| Parameters | <p>same-preference-level — In case that the tunnel-spec selection algorithm evaluates into a tunnel that is currently unavailable (for example tunnel in a blacklist) then the next elected tunnel, if available, will be chosen within the same preference-level as the last attempted tunnel. Only when all tunnels within the same preference level are exhausted, the tunnel selection algorithm will move to the next preference level.</p> <p>In case that a new session setup request is received while all tunnels on the same preference level are blacklisted, the L2TP session will try to be established on blacklisted tunnels before the tunnel selection moves to the next preference level.</p> <p>next-preference-level — In case that the tunnel-spec selection algorithm evaluates into a tunnel that is currently unavailable (for example tunnel in a blacklist) then the selection algorithm will try to select the tunnel from the next preference level, even though the tunnels on the same preference level might be available for selection.</p> <p>Default next-preference-level</p> |
|-------------------|--|

replace-result-code

| | | | |
|--------------------|--|---------------|--|
| Syntax | replace-result-code <i>code</i> [<i>code...</i> (upto 3 max)] no replace-result-code | | |
| Context | config>router>l2tp config>service>vpn>l2tp | | |
| Description | This command will replace CDN Result-Code 4, 5 and 6 on LNS with the Result Code 2. This is needed for interoperability with some implementation of LAC which only take action based on CDN Result-Code 2, while ignore CDN Result-Code 4, 5 and 6. | | |
| Default | no replace-result-code | | |
| Parameters | <i>code</i> — Specifies the L2TP Result codes that need to be replaced. <table> <tr> <td>Values</td><td> cdn-tmp-no-facilities — CDN Result-Code 4 on LNS will be replaced with the result code 2 before it is sent to LAC. cdn-prem-no-facilities — CDN Result-Code 5 on LNS will be replaced with the result code 2 before it is sent to LAC. cdn-inv-dest — CDN Result-Code 6 on LNS will be replaced with the result code 2 before it is sent to LAC. </td></tr> </table> | Values | cdn-tmp-no-facilities — CDN Result-Code 4 on LNS will be replaced with the result code 2 before it is sent to LAC. cdn-prem-no-facilities — CDN Result-Code 5 on LNS will be replaced with the result code 2 before it is sent to LAC. cdn-inv-dest — CDN Result-Code 6 on LNS will be replaced with the result code 2 before it is sent to LAC. |
| Values | cdn-tmp-no-facilities — CDN Result-Code 4 on LNS will be replaced with the result code 2 before it is sent to LAC. cdn-prem-no-facilities — CDN Result-Code 5 on LNS will be replaced with the result code 2 before it is sent to LAC. cdn-inv-dest — CDN Result-Code 6 on LNS will be replaced with the result code 2 before it is sent to LAC. | | |

tunnel-selection-blacklist

| | |
|--------------------|---|
| Syntax | tunnel-selection-blacklist |
| Context | config>router>l2tp |
| Description | This command enables the context to configure L2TP Tunnel Selection Blacklist parameters. |

add-tunnel

| | |
|--------------------|--|
| Syntax | add-tunnel never add-tunnel on <i>reason</i> [<i>reason...</i> (upto 8 max)] no add-tunnel |
| Context | config>router>l2tp>tunnel-selection-blacklist config>service>vpn>l2tp>tunnel-selection-blacklist |
| Description | This command will force the tunnel to the blacklist and render it unavailable for new sessions for the duration of preconfigured time. Peers are always forced to the black list in case that they time out (failure to receive response to control packets). In addition to time outs, certain events can be used to trigger placement of the tunnel on the black list. |
| Default | add-tunnel never |
| Parameters | <i>reason</i> — Specifies the return codes or events that determine which tunnels are added to the blacklist. |

Table 8 **Return codes**

| Return code | Tunnels added to blacklist |
|---------------------------------------|---|
| cdn-err-code | A tunnel will be forced to the blacklist in case that CDN message with the Result Code 2 (Call disconnected for the reasons indicated in error code) is received. |
| cdn-inv-dest | A tunnel will be forced to the blacklist in case that CDN message with the Result Codes 6 (Invalid destination) is received. |
| cdn-tmp-no-facilities | A tunnel will be forced to the blacklist in case that CDN message with the Result Code 4 is received (Call failed due to lack of appropriate facilities being available - temporary condition) is received. |
| cdn-perm-no-facilities | A tunnel will be forced to the blacklist in case that CDN message with the Result Codes 5 (Call failed due to lack of appropriate facilities being available - permanent condition) is received. |
| tx-cdn-not-established-in-time | A tunnel will be forced to the blacklist in case that CDN message with the Result Code 10 (Call was not established within time allotted by LAC) is sent from the LAC to the LNS. |
| stop-ccn-err-code | A tunnel will be forced to the blacklist in case that StopCCN message with the Result Code 2 (General error – Error Code indicates the problem) is sent or received. |

Table 8 Return codes (Continued)

| Return code | Tunnels added to blacklist |
|----------------------------|---|
| stop-ccn-other | <p>A tunnel will be forced to the blacklist in case that StopCCN message with the following Result Codes is received:</p> <ul style="list-style-type: none"> (1) General request to clear control connection (4) Requester is not authorized to establish a control channel (5) Protocol version not supported (6) Requester is being shutdown <p>Or in the case that the StopCCN with the following result codes is transmitted:</p> <ul style="list-style-type: none"> (4) Requester is not authorized to establish a control channel. (5) Protocol version not supported <p>The receipt of the following Result Codes will NEVER blacklist a tunnel:</p> <ul style="list-style-type: none"> (0) Reserved (3) Control channel already exist (7) Finite state machine error (8) Undefined <p>Transmission of the following Result Codes will NEVER blacklist a tunnel:</p> <ul style="list-style-type: none"> (1) General request to clear control connection (3) Control channel already exist (6) Requester is being shutdown (7) Finite state machine error |
| addr-change-timeout | <p>A timed-out tunnel for which the peer IP address has changed mid-session (from the one that is provided initially during configuration) will be forced to the blacklist. In absence of this configuration option, only the configured peer for the tunnel will be blacklisted, but not the tunnel itself which now has a different peer address than the one initially configured.</p> |

never — When specified, no tunnels will be placed on blacklist under any circumstance.
This parameter will available to preserve backward compatibility.

max-list-length

Syntax **max-list-length unlimited**
 max-list-length count
 no max-list-length

| | |
|--------------------|--|
| Context | config>router>l2tp>tunnel-selection-blacklist config>service>vprn>l2tp>tunnel-selection-blacklist |
| Description | <p>This command configured the maximum length of the peer/tunnel blacklist.</p> <p>This command specifies how many items (tunnels or peers) can be in the tunnel-selection-blacklist. If a tunnel or peer needs to be added to the tunnel-selection-blacklist and the tunnel-selection-blacklist is full, the system will remove the item (tunnel or peer) from the blacklist that was in this blacklist for the longest time.</p> |
| Default | max-list-length unlimited |
| Parameters | <p>unlimited — Specifies there is no limit.</p> <p>count — Specifies how many items (tunnels or peers) can be in the tunnel-selection-blacklist.</p> <p>Values 1 to 65635</p> |

max-time

| | |
|--------------------|---|
| Syntax | max-time <i>minutes</i> no max-time |
| Context | config>router>l2tp>tunnel-selection-blacklist config>service>vprn>l2tp>tunnel-selection-blacklist |
| Description | This command configures time for which an entity (peer or a tunnel) are kept in the blacklist. |
| Default | max-time 5 |
| Parameters | <p><i>minutes</i> — Specifies the maximum time a tunnel or peer may remain in the blacklist.</p> <p>Values 1 to 60</p> |

timeout-action

| | |
|--------------------|---|
| Syntax | timeout-action <i>action</i> no timeout-action |
| Context | config>router>l2tp>tunnel-selection-blacklist config>service>vprn>l2tp>tunnel-selection-blacklist |
| Description | This command defines an action that will be executed on the entity (peer/tunnel) in the blacklist once the entity becomes eligible for selection again. |
| Default | timeout-action remove-from-blacklist |

Parameters *action* — Specifies the Action to be taken when a tunnel or peer has been in the blacklist for the max-period of time.

Values *remove-from-blacklist* — The peer or tunnel in the blacklist will be removed completely from the blacklist and made eligible for the selection process once the max-time expires. In this mode of operation, multiple new sessions can be mapped into the same, newly released tunnel from the blacklist. The first such session will try to setup the tunnel, while the other will be buffered until the tunnel establishment process is completed. In case that the tunnel remains unavailable, it will be placed in the blacklist again. Consequently all new sessions will have to be renegotiated over an alternate tunnel.

try-one-session — Once the max-time expired, the peer or tunnel in the blacklist is made available for selection only to a single new session request. Only upon successful tunnel establishment will the incoming new sessions be eligible to be mapped into this tunnel. This behavior will avoid session establishment delays in case that the tunnel just removed from the blacklist is still unavailable.

peer-address-change-policy

Syntax **peer-address-change-policy {accept | ignore | reject}**

Context config>router>l2tp

Description This command specifies what to do in case the system receives a L2TP response from another address than the one the request was sent to.

Default peer-address-change-policy reject

Parameters **accept** — Specifies that this system accepts any source IP address change of received L2TP control messages related to a locally originated tunnel in the state waitReply and rejects any peer address change for other tunnels; in case the new peer IP address is accepted, it is learned and used as destination address in subsequent L2TP messages.

ignore — Specifies that this system ignores any source IP address change of received L2TP control messages, does not learn any new peer IP address and does not change the destination address in subsequent L2TP messages.

reject — Specifies that this system rejects any source IP address change of received L2TP control messages and drops those messages.

receive-window-size

Syntax **receive-window-size [4 to 1024]**
 no receive-window-size

| | |
|--------------------|---|
| Context | config>router>l2tp |
| Description | This command configures the L2TP receive window size. |
| Default | receive-window-size 64 |

rtm-debounce-time

| | |
|--------------------|--|
| Syntax | rtm-debounce-time <i>debounce-time</i> no rtm-debounce-time |
| Context | config>router>l2tp |
| Description | <p>This command configures the amount of time, in milliseconds, that the system will wait before declaring an L2TP tunnel down when the remote endpoint IP address cannot be resolved to an active IP route in the local routing table.</p> <p>The default behavior is for the L2TP tunnel to not be declared down based on the remote endpoint IP address reachability.</p> <p>The no form of this command returns the rtm-debounce-time to a value of 0.</p> |
| Default | no rtm-debounce-time |
| Parameters | <i>debounce-time</i> — Specifies the amount of time, in milliseconds, that the system will wait before declaring the associated L2TP tunnel as down. |
| Values | 0 to 5000 |

group

| | |
|--------------------|--|
| Syntax | group <i>tunnel-group-name</i> [create] no group <i>tunnel-group-name</i> |
| Context | config>router>l2tp |
| Description | This command configures an L2TP tunnel group. |
| Parameters | <i>tunnel-group-name</i> — Specifies a name string to identify a L2TP group up to 63 characters in length. create — Mandatory keyword when creating a tunnel group name. The create keyword requirement can be enabled or disabled in the environment>create context. |

session-limit

| | |
|---------------|---|
| Syntax | session-limit <i>session-limit</i> |
|---------------|---|

| | |
|--------------------|---|
| | session-limit unlimited no session-limit |
| Context | config>router>l2tp |
| Description | This command configures the L2TP session limit for the router. L2TP is connection-oriented. The L2TP Network Server (LNS) and LAC maintain state for each call that is initiated or answered by an LAC. An L2TP session is created between the LAC and LNS when an end-to-end PPP connection is established between a remote system and the LNS. Datagrams related to the PPP connection are sent over the tunnel between the LAC and LNS. There is a one to one relationship between established L2TP sessions and their associated calls. |
| Default | no session-limit |
| Parameters | <i>session-limit</i> — Specifies the number of sessions allowed. Values 1 to 131071 unlimited — Specifies to use the maximum number of sessions available. |

avp-hiding

| | |
|--------------------|--|
| Syntax | avp-hiding <i>sensitive</i> <i>always</i> no avp-hiding |
| Context | config>router>l2tp>group |
| Description | This command configures Attribute Value Pair (AVP) hiding. This capability can be used to avoid the passing of sensitive data, such as user passwords, as cleartext in an AVP. The no form of the command returns the value to never allow AVP hiding. |
| Default | no avp-hiding |
| Parameters | <i>avp-hiding</i> — Specifies the method to be used for the authentication of the tunnels in this L2TP group. Default no avp-hiding Values <i>sensitive</i> — AVP hiding is used only for sensitive information (such as username/password). <i>always</i> — AVP hiding is always used. |

challenge

| | |
|----------------|---|
| Syntax | challenge <i>always</i> no challenge |
| Context | config>router>l2tp>group |

| | |
|--------------------|---|
| Description | This command configures the use of challenge-response authentication. The no form of the command reverts to the default never value. |
| Default | no challenge |
| Parameters | <i>always</i> — Specifies that the challenge-response authentication is always used. |
| Default | no challenge |
| Values | always |

df-bit-lac

| | |
|--------------------|---|
| Syntax | df-bit-lac {always never} no df-bit-lac |
| Context | config>router>l2tp config>service>vpn>l2tp |
| Description | By default, the LAC df-bit-lac is always set and sends all L2TP packets with the DF bit set to 1. The DF bit is configurable to allow downstream routers to fragment the L2TP packets. The LAC itself will not fragment L2TP packets. L2TP packets that have a larger MTU size than what the LAC egress ports allows are dropped. |
| Default | df-bit-lac always |
| Parameters | always — Specifies that the LAC will send all L2TP packets with the DF bit set to 1. never — Specifies that the LAC will send all L2TP packets with the DF bit set to 0. |

df-bit-lac

| | |
|--------------------|---|
| Syntax | df-bit-lac {always never default} no df-bit-lac |
| Context | config>router/service>vpn>l2tp>group config>router/service>vpn>l2tp>group>tunnel |
| Description | By default, the LAC df-bit-lac is set to default and sends all L2TP packets with the DF bit set to 1. The DF bit is configurable to allow downstream routers to fragment the L2TP packets. The LAC itself will not fragment L2TP packets. L2TP packets that have a larger MTU size than what the LAC egress ports allows are dropped. The configuration of the df-bit can be overridden at different levels: l2tp, tunnel, and group. The configuration at the tunnel level overrides the configuration on both group and l2tp. The configuration at the group level overrides the configuration on l2tp. |
| Default | df-bit-lac default |
| Parameters | always — Specifies that the LAC will send all L2TP packets with the DF bit set to 1. |

never — Specifies that the LAC will send all L2TP packets with the DF bit set to 0.

default — Follows the DF-bit configuration specified on upper levels.

destruct-timeout

| | |
|--------------------|---|
| Syntax | destruct-timeout <i>destruct-timeout</i> no destruct-timeout |
| Context | config>router>l2tp>group config>router>l2tp>group>tunnel |
| Description | This command configures the period of time that the data of a disconnected tunnel will persist before being removed. The no form of the command removes the value from the configuration. |
| Default | no destruct-timeout |
| Parameters | <i>destruct-timeout</i> — Specifies the automatic removal of dynamic L2TP sessions, in seconds, that are no longer active. Default no destruct-timeout Values 60 to 86400 |

hello-interval

| | |
|--------------------|--|
| Syntax | hello-interval <i>hello-interval</i> no hello-interval |
| Context | config>router>l2tp>group |
| Description | This command configures the time interval between two consecutive tunnel Hello messages. The Hello message is an L2TP control message sent by either peer of a LAC-LNS control connection. This control message is used as a keepalive for the tunnel. The no form of the command removes the interval from the configuration. |
| Default | no hello-interval |
| Parameters | <i>hello-interval</i> — Specifies the time interval, in seconds, between two consecutive tunnel Hello messages. Default no hello-interval Values 60 to 3600 |

idle-timeout

| | |
|--------------------|--|
| Syntax | idle-timeout <i>idle-timeout</i> no idle-timeout |
| Context | config>router>l2tp>group |
| Description | <p>This command configures the period of time that an established tunnel with no active sessions will persist before being disconnected.</p> <p>Enter the no form of the command to maintain a persistent tunnel.</p> <p>The no form of the command removes the idle timeout from the configuration.</p> |
| Default | no idle-timeout |
| Parameters | <i>idle-timeout</i> — Specifies the idle timeout value, in seconds until the group is removed. Default no idle-timeout Values 0 to 3600 |

lns-group

| | |
|--------------------|---|
| Syntax | lns-group <i>lns-group-id</i> no lns-group |
| Context | config>router>l2tp>group |
| Description | This command configures the ISA LNS group. |
| Default | no lns-group |
| Parameters | <i>lns-group-id</i> — Specifies the LNS group ID. Values 1 to 4 |

load-balance-method

| | |
|--------------------|--|
| Syntax | load-balance-method { per-session per-tunnel } no load-balance-method |
| Context | config>router>l2tp>group config>router>l2tp>group>tunnel |
| Description | This command describes how new sessions are assigned to an L2TP ISA MDA. |
| Default | load-balance-method per-session |

| | |
|-------------------|---|
| Parameters | <p>per-session — Specifies that the lowest granularity for load-balancing is a session; each session can be assigned to a different ISA MDA.</p> <p>per-tunnel — Specifies that the lowest granularity for load-balancing is a tunnel; all sessions associated with the same tunnel are assigned to the same ISA MDA; this may be useful or required in certain cases, for example:</p> <ul style="list-style-type: none"> MLPPP with multiple links per bundle; HPol intermediate destination arbiters where the intermediate destination is an L2TP tunnel. |
|-------------------|---|

local-address

| | |
|--------------------|---|
| Syntax | <p>local-address <i>ip-address</i></p> <p>no local-address</p> |
| Context | config>router>l2tp>group>tunnel |
| Description | This command configures the local address. |
| Default | no local-address |
| Parameters | <i>ip-address</i> — Specifies the IP address used during L2TP authentication. |

local-name

| | |
|--------------------|--|
| Syntax | <p>local-name <i>host-name</i></p> <p>no local-name</p> |
| Context | <p>config>router>l2tp>group</p> <p>config>router>l2tp>group>tunnel</p> |
| Description | <p>This command creates the local host name used by this system for the tunnels in this L2TP group during the authentication phase of tunnel establishment. It can be used to distinguish tunnels.</p> <p>The no form of the command removes the name from the configuration.</p> |
| Default | no local-name |
| Parameters | <p><i>host-name</i> — Specifies the host name, up to 64 characters in length, that the router will use to identify itself during L2TP authentication.</p> <p>Default no local-name</p> |

max-retries-estab

| | |
|--------------------|---|
| Syntax | max-retries-estab <i>max-retries</i> no max-retries-estab |
| Context | config>router>l2tp>group config>router>l2tp>group>tunnel |
| Description | This command configures the number of retries allowed for this L2TP tunnel while it is established, before its control connection goes down. The no form of the command removes the value from the configuration. |
| Default | no max-retries-estab |
| Parameters | <i>max-retries</i> — Specifies the maximum number of retries for an established tunnel. Default no max-retries-estab Values 2 to 7 |

max-retries-not-estab

| | |
|--------------------|---|
| Syntax | max-retries-not-estab <i>max-retries</i> no max-retries-not-estab |
| Context | config>router>l2tp>group config>router>l2tp>group>tunnel |
| Description | This command configures the number of retries allowed for this L2TP tunnel while it is not established, before its control connection goes down. The no form of the command removes the value from the configuration. |
| Default | no max-retries-not-estab |
| Parameters | <i>max-retries</i> — Specifies the maximum number of retries for non-established tunnels. Default no max-retries-not-estab Values 2 to 7 |

password

| | |
|----------------|--|
| Syntax | password <i>password</i> [hash hash2] no password |
| Context | config>router>l2tp>group config>router>l2tp>group>tunnel config>router>l2tp>l2tpv3 |


```
config>router>l2tp>group>l2tpv3
```

Description This command configures the password between L2TP LAC and LNS

The **no** form of the command removes the password.

Default no password

Parameters *password* — Configures the password used for challenge/response calculation and AVP hiding. The maximum length can be up to 20 characters if unhashed, 32 characters if hashed, 54 characters if the hash2 keyword is specified.

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

ppp

Syntax **ppp**

Context config>router>l2tp>group
config>router>l2tp>group>tunnel

Description This command configures PPP for the L2TP tunnel group.

authentication

Syntax **authentication {chap | pap | pref-chap | prep-pap}**

Context config>router>l2tp>group>ppp

Description This command configures the PPP authentication protocol to negotiate authentication.

Default authentication pref-chap

Parameters **chap** — Specifies to always use CHAP for authentication.

pap — Specifies to always use PAP for authentication.

pref-chap — Specifies to use CHAP as the preferred authentication method, and to use PAP if that attempt fails.

pref-pap — Specifies to use PAP as the preferred authentication method, and to use CHAP if that attempt fails.

authentication-policy

| | |
|--------------------|---|
| Syntax | authentication-policy <i>auth-policy-name</i> no authentication-policy |
| Context | config>router>l2tp>group>ppp |
| Description | This command configures the authentication policy. |
| Default | no authentication-policy |
| Parameters | <i>auth-policy-name</i> — Specifies the authentication policy name. Values 32 chars max |

default-group-interface

| | |
|--------------------|---|
| Syntax | default-group-interface <i>ip-int-name</i> service-id <i>service-id</i> no default-group-interface |
| Context | config>router>l2tp>group>ppp |
| Description | This command configures the default group interface. |
| Default | no default-group-interface |
| Parameters | <i>ip-int-name</i> — Specifies the interface name. Values 32 chars max <i>service-id</i> — Specifies the service ID. Values 1 to 2147483648 <i>svc-name</i> — Specifies the service name (instead of service ID). Values 64 chars max |

keepalive

| | |
|--------------------|--|
| Syntax | keepalive <i>seconds</i> [hold-up-multiplier <i>multiplier</i>] no keepalive |
| Context | config>router>l2tp>group>ppp |
| Description | This command configures the PPP keepalive interval and multiplier. |
| Default | keepalive 30 hold-up-multiplier 3 |
| Parameters | <i>seconds</i> — Specifies in seconds the interval. Values 10 to 300 |

multiplier — Specifies the multiplier.

Values 1 to 5

lcp-force-ack-accm

| | |
|--------------------|--|
| Syntax | [no] lcp-force-ack-accm |
| Context | config>router>l2tp>group>ppp config>router>l2tp>group>tunnel>ppp |
| Description | This command enables or disables the LCP Asynchronous Control Character Map (ACCM) configuration option. When the ACCM configuration option is enabled, the option is acknowledged during the LCP negotiation between the LNS and the PPP client, but no ACCM mapping is performed. By default, the ACCM configuration option is rejected. The no form of this command reverts to the default value. |
| Default | no lcp-force-ack-accm |

mtu

| | |
|--------------------|--|
| Syntax | mtu <i>mtu-bytes</i> no mtu |
| Context | config>router>l2tp>group>ppp |
| Description | This command configures the maximum PPP MTU size. |
| Default | mtu 1500 |
| Parameters | <i>mtu-bytes</i> — Specifies, in bytes, the maximum PPP MTU size. Values 512 to 9212 |

proxy-authentication

| | |
|--------------------|---|
| Syntax | [no] proxy-authentication |
| Context | config>router>l2tp>group>ppp |
| Description | This command configures the use of the authentication AVPs received from the LAC. |
| Default | no proxy-authentication |

proxy-lcp

| | |
|--------------------|--|
| Syntax | [no] proxy-lcp |
| Context | config>router>l2tp>group>ppp |
| Description | This command configures the use of the proxy LCP AVPs received from the LAC. |
| Default | no proxy-lcp |

user-db

| | |
|--------------------|---|
| Syntax | user-db <i>local-user-db-name</i> no user-db |
| Context | config>router>l2tp>group>ppp |
| Description | This command configures the local user database to use for PPP PAP/CHAP authentication. |
| Default | no user-db |
| Parameters | <i>local-user-db-name</i> — Specifies the local user database name. Values 32 chars max |

session-assign-method

| | |
|--------------------|--|
| Syntax | session-assign-method [existing-first weighted weighted-random] no session-assign-method |
| Context | config>router>l2tp>group |
| Description | This command specifies how new sessions are assigned to one of the set of suitable tunnels that are available or could be made available. |
| Default | existing-first |
| Parameters | existing-first — Specifies that all new sessions are placed by preference in the existing tunnels. weighted — Specifies that the sessions are shared between the available tunnels. If necessary, new tunnels are set up until the maximum number is reached. The distribution aims at an equal ratio of the actual number of sessions to the maximum number of sessions. weighted-random — Enhances the weighted algorithm such that when there are multiple tunnels with an equal number of sessions (equal weight), LAC randomly selects a tunnel. |

session-limit

| | |
|--------------------|--|
| Syntax | session-limit <i>session-limit</i> session-limit unlimited no session-limit |
| Context | config>router>l2tp>group config>router>l2tp>group>tunnel |
| Description | This command configures the session limit. The value controls how many L2TP session will be allowed within a given context (system, group, tunnel). The no form of the command removes the value from the configuration. |
| Default | no session-limit |
| Parameters | <i>session-limit</i> — Specifies the allowed number of sessions within the given context. Values 1 to 131071 unlimited — Specifies to use the maximum number of sessions available. |

2.12.2.3.1 Router L2TP Tunnel Commands

Router L2TP tunnel commands only apply to the 7750 SR and 7450 ESS.

tunnel

| | |
|--------------------|---|
| Syntax | tunnel <i>tunnel-name</i> [create] no tunnel <i>tunnel-name</i> |
| Context | config>router>l2tp>group |
| Description | This command configures an L2TP tunnel. A tunnel exists between a LAC-LNS pair and consists of a Control Connection and zero or more L2TP sessions. The tunnel carries encapsulated PPP datagrams and control messages between the LAC and the L2TP Network Server (LNS). |
| Parameters | <i>tunnel-name</i> — Specifies a valid string to identify a L2TP up to 32 characters in length. create — Mandatory while creating a new tunnel. |

auto-establish

| | |
|----------------|-------------------------------------|
| Syntax | [no] auto-establish |
| Context | config>router>l2tp>group>tunnel |

| | |
|--------------------|--|
| Description | This command specifies if this tunnel is to be automatically set up by the system. |
| Default | no auto-establish |

avp-hiding

| | |
|--------------------|---|
| Syntax | avp-hiding { never sensitive always } no avp-hiding |
| Context | config>router>l2tp>group>tunnel |
| Description | <p>This command configures Attribute Value Pair (AVP) hiding. This capability can be used to avoid the passing of sensitive data, such as user passwords, as cleartext in an AVP.</p> <p>It is recommended that sensitive information not be sent in clear text.</p> <p>The no form of the command removes the parameter of the configuration and indicates that the value on group level will be taken.</p> |
| Default | no avp-hiding |
| Parameters | <i>avp-hiding</i> — Specifies the method to be used for the authentication of the tunnel. Values never — AVP hiding is not used. sensitive — AVP hiding is used only for sensitive information (such as username/password). always — AVP hiding is always used. |

challenge

| | |
|--------------------|--|
| Syntax | challenge <i>challenge-mode</i> no challenge |
| Context | config>router>l2tp>group>tunnel |
| Description | <p>This command configures the use of challenge-response authentication.</p> <p>The no form of the command removes the parameter from the configuration and indicates that the value on group level will be taken.</p> |
| Default | no challenge |
| Parameters | <i>challenge-mode</i> — Specifies when challenge-response is to be used for the authentication of the tunnel. Values always — Always allows the use of challenge-response authentication. never — Never allows the use of challenge-response authentication. |

hello-interval

| | |
|--------------------|--|
| Syntax | hello-interval <i>hello-interval</i> hello-interval infinite no hello-interval |
| Context | config>router>l2tp>group>tunnel |
| Description | This command configures the number of seconds between sending Hellos for a L2TP tunnel. The no form removes the parameter from the configuration and indicates that the value on group level will be taken. |
| Default | no hello-interval |
| Parameters | <i>hello-interval</i> — Specifies the time interval, in seconds, between two consecutive tunnel Hello messages. Values 60 to 3600 infinite — Specifies that no hello messages are sent. |

idle-timeout

| | |
|--------------------|---|
| Syntax | idle-timeout <i>idle-timeout</i> idle-timeout infinite no idle-timeout |
| Context | config>router>l2tp>group>tunnel |
| Description | This command configures the idle timeout to wait before being disconnect. The no form indicates that the parameter will be removed from the configuration and that the value specified on group level will be taken. |
| Default | no idle-timeout |
| Parameters | <i>idle-timeout</i> — Specifies the idle timeout, in seconds. Values 0 to 3600 infinite — Specifies that the tunnel will not be closed when idle. |

peer

| | |
|--------------------|---|
| Syntax | peer <i>ip-address</i> no peer |
| Context | config>router>l2tp>group>tunnel |
| Description | This command configures the peer address. |

The **no** form of the command removes the IP address from the tunnel configuration.

| | |
|-------------------|---|
| Default | no peer |
| Parameters | <i>ip-address</i> — Sets the LNS IP address for the tunnel. |

preference

| | |
|--------------------|--|
| Syntax | preference <i>preference</i> no preference |
| Context | config>router>l2tp>group>tunnel |
| Description | This command configures a preference number that indicates the relative preference assigned to a tunnel when using a weighted session assignment. The no form of the command removes the preference value from the tunnel configuration. |
| Default | no preference |
| Parameters | <i>preference</i> — Specifies the tunnel preference number with its group. The value 0 corresponds to the highest preference. Values 0 to 16777215 |

remote-name

| | |
|--------------------|---|
| Syntax | remote-name <i>host-name</i> no remote-name |
| Context | config>router>l2tp>group>tunnel |
| Description | This command configures a string to be compared to the host name used by the tunnel peer during the authentication phase of tunnel establishment. |
| Default | no remote-name |
| Parameters | <i>host-name</i> — Specifies a remote host name for the tunnel up to 64 characters in length. |

tunnel-selection-blacklist

| | |
|--------------------|---|
| Syntax | tunnel-selection-blacklist |
| Context | config>router>l2tp |
| Description | This command enables the context to configure L2TP Tunnel Selection Blacklist parameters. |

add-tunnel

| | |
|---------------------------------------|--|
| Syntax | add-tunnel never add-tunnel on <i>reason</i> [<i>reason...</i> (upto 8 max)] no add-tunnel |
| Context | config>router>l2tp>tunnel-selection-blacklist config>service>vpn>l2tp>tunnel-selection-blacklist |
| Description | This command will force the tunnel to the blacklist and render it unavailable for new sessions for the duration of preconfigured time. Peers are always forced to the black list in case that they time out (failure to receive response to control packets). In addition to time outs, certain events can be used to trigger placement of the tunnel on the black list. |
| Default | add-tunnel never |
| Parameters | <i>reason</i> — Specifies the return codes or events that determine which tunnels are added to the blacklist. |
| cdn-err-code | A tunnel will be forced to the blacklist in case that CDN message with the Result Code 2 (Call disconnected for the reasons indicated in error code) is received. |
| cdn-inv-dest | A tunnel will be forced to the blacklist in case that CDN message with the Result Codes 6 (Invalid destination) is received. |
| cdn-tmp-no-facilities | A tunnel will be forced to the blacklist in case that CDN message with the Result Code 4 is received (Call failed due to lack of appropriate facilities being available - temporary condition) is received. |
| cdn-perm-no-facilities | A tunnel will be forced to the blacklist in case that CDN message with the Result Codes 5 (Call failed due to lack of appropriate facilities being available - permanent condition) is received. |
| tx-cdn-not-established-in-time | A tunnel will be forced to the blacklist in case that CDN message with the Result Code 10 (Call was not established within time allotted by LAC) is sent from the LAC to the LNS. |
| stop-ccn-err-code | A tunnel will be forced to the blacklist in case that StopCCN message with the Result Code 2 (General error – Error Code indicates the problem) is sent or received. |

| | |
|----------------------------|---|
| stop-ccn-other | <p>A tunnel will be forced to the blacklist in case that StopCCN message with the following Result Codes is received:</p> <ul style="list-style-type: none"> (1) General request to clear control connection (4) Requester is not authorized to establish a control channel (5) Protocol version not supported (6) Requester is being shutdown <p>Or in the case that the StopCCN with the following result codes is transmitted:</p> <ul style="list-style-type: none"> (4) Requester is not authorized to establish a control channel. (5) Protocol version not supported <p>The receipt of the following Result Codes will NEVER blacklist a tunnel:</p> <ul style="list-style-type: none"> (0) Reserved (3) Control channel already exist (7) Finite state machine error (8) Undefined <p>Transmission of the following Result Codes will NEVER blacklist a tunnel:</p> <ul style="list-style-type: none"> (1) General request to clear control connection (3) Control channel already exist (6) Requester is being shutdown (7) Finite state machine error |
| addr-change-timeout | <p>A timed-out tunnel for which the peer IP address has changed mid-session (from the one that is provided initially during configuration) will be forced to the blacklist. In absence of this configuration option, only the configured peer for the tunnel will be blacklisted, but not the tunnel itself which now has a different peer address than the one initially configured.</p> |

never — When specified, no tunnels will be placed on blacklist under any circumstance.
This parameter will be available to preserve backward compatibility.

max-list-length

| | |
|--------------------|---|
| Syntax | <p>max-list-length unlimited max-list-length <i>count</i> no max-list-length</p> |
| Context | <p>config>router>l2tp>tunnel-selection-blacklist config>service>vpn>l2tp>tunnel-selection-blacklist</p> |
| Description | <p>This command configured the maximum length of the peer/tunnel blacklist.</p> |

This command specifies how many items (tunnels or peers) can be in the tunnel-selection-blacklist. If a tunnel or peer needs to be added to the tunnel-selection-blacklist and the tunnel-selection-blacklist is full, the system will remove the item (tunnel or peer) from the blacklist that was in this blacklist for the longest time.

| | |
|-------------------|---|
| Default | max-list-length unlimited |
| Parameters | unlimited — Specifies there is no limit. count — Specifies how many items (tunnels or peers) can be in the tunnel-selection-blacklist. Values 1 to 65635 |

max-time

| | |
|--------------------|--|
| Syntax | max-time <i>minutes</i> no max-time |
| Context | config>router>l2tp>tunnel-selection-blacklist config>service>vpn>l2tp>tunnel-selection-blacklist |
| Description | This command configures time for which an entity (peer or a tunnel) are kept in the blacklist. |
| Default | max-time 5 |
| Parameters | <i>minutes</i> — Specifies the maximum time a tunnel or peer may remain in the blacklist. Values 1 to 60 |

timeout-action

| | |
|--------------------|---|
| Syntax | timeout-action <i>action</i> no timeout-action |
| Context | config>router>l2tp>tunnel-selection-blacklist config>service>vpn>l2tp>tunnel-selection-blacklist |
| Description | This command defines an action that will be executed on the entity (peer/tunnel) in the blacklist once the entity becomes eligible for selection again. |
| Default | timeout-action remove-from-blacklist |

Parameters *action* — Specifies the action to be taken when a tunnel or peer has been in the blacklist for the maximum period of time.

Values **remove-from-blacklist** — The peer or tunnel in the blacklist will be removed completely from the blacklist and made eligible for the selection process once the maximum period expires. In this mode of operation, multiple new sessions can be mapped into the same, newly-released tunnel from the blacklist. The first such session will try to set up the tunnel, while the other will be buffered until the tunnel establishment process is completed. If the tunnel remains unavailable, it will be placed in the blacklist again. Consequently, all new sessions will have to be renegotiated over an alternate tunnel.

try-one-session — Once the maximum period expires, the peer or tunnel in the blacklist is made available for selection only to a single new session request. Only upon successful tunnel establishment will the incoming new sessions be eligible to be mapped into this tunnel. This behavior will avoid session establishment delays if the tunnel just removed from the blacklist is still unavailable.

2.12.2.4 Router Interface Commands

interface

Syntax **[no] interface** *ip-int-name* [**unnumbered-mpls-tp** | **gmpls-loopback** | **control-tunnel**]

Context config>router

Description This command creates a logical IP routing or unnumbered MPLS-TP interface. Once created, attributes like IP address, port, or system can be associated with the IP interface.

Interface names are case-sensitive and must be unique within the group of IP interfaces defined for **config router interface** and **config service ies interface**. Interface names must not be in the dotted decimal notation of an IP address.; for example, the name “1.1.1.1” is not allowed, but “int-1.1.1.1” is allowed. Show commands for router interfaces use either the interface names or the IP addresses. Ambiguity can exist if an IP address is used as an IP address and an interface name. Duplicate interface names can exist in different router instances, although this is not recommended because it is confusing.

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

Although not a keyword, the ip-int-name “**system**” is associated with the network entity (such as a specific router), not a specific interface. The system interface is also referred to as the loopback address.

An unnumbered MPLS-TP interface is a special type of interface that is only intended for MPLS-TP LSPs. IP routing protocols are blocked on interfaces of this type. If an interface is configured as **unnumbered-mpls-tp**, then it can only be associated with an Ethernet port or VLAN, using the **port** command, then either a unicast, multicast, or broadcast remote MAC address may be configured. Only static ARP is supported.

A GMPLS loopback interface is a special type of loopback interface that is used as the IP interface for a GMPLS IP Control Channel (IPCC). RSVP and LMP packets associated with GMPLS are associated with this loopback interface. All other IP protocols are blocked on this interface. One **gmpls-loopback** interface is required for each GMPLS peer node.

The **control-tunnel** parameter creates a loopback interface representing a GRE tunnel. One IP tunnel can be created in this interface.

Only the primary IPv4 interface address and only one IP tunnel per interface are allowed. Multiple tunnels can be configured using up to four **controlTunnel** loopback interfaces. A static route can take the new **controlTunnel** interface as a next hop.

The **no** form of the command removes the IP interface and all the associated configurations. The interface must be administratively shut down before issuing the **no interface** command.

| | |
|-------------------|--|
| Default | No interfaces or names are defined within the system. |
| Parameters | <p><i>ip-int-name</i> — The name of the IP interface. Interface names must be unique within the group of defined IP interfaces for config router interface and config service ies interface commands. An interface name cannot be in the form of an IP address. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes.</p> <p>Values 1 to 32 alphanumeric characters.</p> <p>If the <i>ip-int-name</i> already exists, the context is changed to maintain that IP interface. If <i>ip-int-name</i> already exists within another service ID or is an IP interface defined within the config router commands, an error will occur and the context will not be changed to that IP interface. If <i>ip-int-name</i> does not exist, the interface is created and the context is changed to that interface for further command processing.</p> <p>unnumbered-mpls-tp — Specifies that an interface is an unnumbered MPLS-TP. An unnumbered MPLS-TP interface is a special type of interface that is only intended for MPLS-TP LSPs. IP routing protocols are blocked on interfaces of this type. If an interface is configured as unnumbered-mpls-tp, then it can only be associated with an Ethernet port or VLAN, using the port command. A unicast, multicast, or broadcast remote MAC address can be configured using the static-arp command. Only static ARP is supported.</p> <p>gmpls-loopback — Specifies that the interface is a loopback interface for GMPLS control plane packets.</p> <p>control-tunnel — Specifies that the interface represents a loopback interface for a GRE tunnel to be used for the GMPLS IPCC.</p> |

address

| | |
|--------------------|--|
| Syntax | address { <i>ip-address/mask</i> ip-address <i>netmask</i> } [broadcast all-ones host-ones] [track-srrp <i>srrp-instance</i>] no address |
| Context | config>router>if |
| Description | <p>This command assigns an IP address, IP subnet, and broadcast address format to an IP interface. Only one IP address can be associated with an IP interface. Use the secondary command to assign additional addresses.</p> <p>An IP address must be assigned to each IP interface. An IP address and a mask combine to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the router.</p> <p>The local subnet that the address command defines must not be part of the services address space within the routing context by use of the config router service-prefix command. Once a portion of the address space is allocated as a service prefix, that portion is not available to IP interfaces for network core connectivity.</p> <p>The IP address for the interface can be entered in either CIDR (Classless Inter-Domain Routing) or traditional dotted decimal notation. Show commands display CIDR notation and are stored in configuration files.</p> <p>By default, no IP address or subnet association exists on an IP interface until it is explicitly created.</p> <p>The no form of the command removes the IP address assignment from the IP interface. Interface specific configurations for MPLS are also removed. This will operationally stop any MPLS LSPs that explicitly reference that IP address. When a new IP address is configured, interface specific configurations for MPLS need to be added. IEEE 1588 port based timestamping configured with ptp-hw-assist is also disabled.</p> |
| Default | No IP address is assigned to the IP interface. |
| Parameters | <p><i>ip-address</i> — The IP address of the IP interface. The <i>ip-addr</i> portion of the address command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.</p> <p>Values 1.0.0.0 to 223.255.255.255</p> <p><i>/</i> — The forward slash is a parameter delimiter that separates the <i>ip-addr</i> portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the <i>ip-addr</i>, the “/” and the <i>mask-length</i> parameter. If a forward slash does not immediately follow the <i>ip-addr</i>, a dotted decimal mask must follow the prefix.</p> |

mask-length — The subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-addr* from the *mask-length* parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address. Allowed values are integers in the range 1— 32. A mask length of 32 is reserved for system IP addresses.

Values 1 to 32

mask — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-addr* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that will be used in a logical 'AND' function to derive the local subnet of the IP address. A mask of 255.255.255.255 is reserved for system IP addresses.

Values 128.0.0.0 to 255.255.255.255

netmask — The subnet mask in dotted decimal notation.

Values 0.0.0.0 to 255.255.255.255 (network bits all 1 and host bits all 0)

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

The **all-ones** keyword following the **broadcast** parameter specifies that the broadcast address used by the IP interface for this IP address will be 255.255.255.255, also known as the local broadcast.

The **host-ones** keyword following the **broadcast** parameter specifies that the broadcast address used by the IP interface for this IP address will be the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the *ip-addr* and the *mask-length* or *mask* with all the host bits set to binary 1. This is the default broadcast address used by an IP interface.

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

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


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

Default host-ones

Values all-ones, host-ones

track-srrp — Specifies the SRRP instance ID that this interface route needs to track.

allow-directed-broadcasts

| | |
|--------------------|---|
| Syntax | [no] allow-directed-broadcasts |
| Context | config>router>if |
| Description | <p>This command enables the forwarding of directed broadcasts out of the IP interface.</p> <p>A directed broadcast is a packet received on a local router interface destined for the subnet broadcast address of another IP interface. The allow-directed-broadcasts command on an IP interface enables or disables the transmission of packets destined to the subnet broadcast address of the egress IP interface.</p> <p>When enabled, a frame destined to the local subnet on this IP interface is sent as a subnet broadcast out this interface.</p> <p> Note: Allowing directed broadcasts is a well-known mechanism used for denial-of-service attacks.</p> <p>By default, directed broadcasts are not allowed and are discarded at this egress IP interface.</p> <p>The no form of the command disables directed broadcasts forwarding out of the IP interface.</p> |
| Default | no allow-directed-broadcasts — Directed broadcasts are dropped. |

arp-limit

| | |
|--------------------|--|
| Syntax | arp-limit <i>limit</i> [log-only] [threshold <i>percent</i>] no arp-limit |
| Context | config>router>if |
| Description | <p>This command configures the maximum amount of dynamic IPv4 ARP entries that can be learned on an IP interface.</p> <p>When the number of dynamic ARP entries reaches the configured percentage of this limit, an SNMP trap is sent. When the limit is exceeded, no new entries are learned until an entry expires and traffic to these destinations will be dropped. Entries that have already been learned will be refreshed.</p> <p>The no form of the command removes the arp-limit.</p> |
| Default | 90 percent |
| Parameters | log-only — Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, entries above the limit will be learned. |

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

Values 0 to 100

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

Values 0 to 524288

arp-timeout

| | |
|--------------------|---|
| Syntax | arp-timeout <i>seconds</i> no arp-timeout |
| Context | config>router>if |
| Description | This command configures the minimum time, in seconds, an ARP entry learned on the IP interface is stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host. Otherwise, the ARP entry is aged from the ARP table. If the arp-timeout value is set to 0 seconds, ARP aging is disabled. The no form of the command reverts to the default value. |
| Default | 14400 seconds (4 hours) |
| Parameters | <i>seconds</i> — The minimum number of seconds a learned ARP entry is stored in the ARP table, expressed as a decimal integer. A value of 0 specifies that the timer is inoperative and learned ARP entries will not be aged. Values 0 to 65535 |

bfd

| | |
|--------------------|--|
| Syntax | bfd <i>transmit-interval</i> [receive <i>receive-interval</i>] [multiplier <i>multiplier</i>] [echo-receive <i>echo-interval</i>] [type <i>cpm-np</i>] no bfd |
| Context | config>router>if config>router>if>ipv6 |
| Description | This command specifies the bidirectional forwarding detection (BFD) parameters for the associated IP interface. If no parameters are defined the default values are used. The multiplier specifies the number of consecutive BFD messages that must be missed from the peer before the BFD session state is changed to down and the upper level protocols (OSPF, IS-IS, BGP or PIM) is notified of the fault. |

The **no** form of the command removes BFD from the router interface regardless of the IGP/RSVP.

Important notes: On the 7750 SR and 7950 XRS SR OS, the *transmit-interval* and **receive receive-interval** values can only be modified to a value less than 100 ms when:

1. The **type cpm-np option** is explicitly configured.
2. The service is shut down (**shutdown**)
3. The interval is specified 10 to 100000.
4. The service is re-enabled (**no shutdown**)

To remove the **type cpm-np** option, re-issue the **bfd** command without specifying the **type** parameter.

| | |
|-------------------|--|
| Default | no bfd |
| Parameters | <p><i>transmit-interval</i> — Sets the transmit interval, in milliseconds, for the BFD session.</p> <p>Values 10 to 100000 (see Important Notes above) The minimum value is 300 msec for central BFD sessions in the 7950 XRS.</p> <p>Default 100</p> <p><i>receive receive-interval</i> — Sets the receive interval, in milliseconds, for the BFD session.</p> <p>Values 10 to 100000 (see Important Notes above)</p> <p>Default 100</p> <p><i>multiplier multiplier</i> — Sets the multiplier for the BFD session.</p> <p>Values 3 to 20</p> <p>Default 3</p> <p><i>echo-receive echo-interval</i> — Sets the minimum echo receive interval, in milliseconds, for the session.</p> <p>Values 100 to 100000</p> <p>Default 0</p> <p>type cpm-np — Selects the CPM network processor as the local termination point for the BFD session for the 7750 SR and 7950 XRS. See Important Notes, above.</p> |

cflowd-parameters

| | |
|----------------|---|
| Syntax | cflowd-parameters no cflowd-parameters |
| Context | config>router>if |

| | |
|--------------------|---|
| Description | <p>This command creates the configuration context to configure cflowd parameters for the associated IP interfaces.</p> <p>cflowd is used for network planning and traffic engineering, capacity planning, security, application and user profiling, performance monitoring, usage-based billing, and SLA measurement.</p> <p>At a minimum, the sampling command must be configured within this context in order to enable cflowd sampling, otherwise traffic sampling will not occur.</p> |
| Default | no cflowd-parameters |

sampling

| | |
|--------------------|---|
| Syntax | <p>sampling {unicast multicast} type {acl interface} [direction {ingress-only egress-only both}]</p> <p>no sampling {unicast multicast}</p> |
| Context | config>router>if>cflowd-parameters |
| Description | <p>This command enables and configures the cflowd sampling behavior to collect traffic flow samples through a router for analysis.</p> <p>This command can be used to configure the sampling parameters for unicast and multicast traffic separately. If sampling is not configured for either unicast or multicast traffic, then that type of traffic will not be sampled.</p> <p>If cflowd is enabled without either egress-only or both specified or with the ingress-only keyword specified, then only ingress sampling will be enabled on the associated IP interface.</p> <p>The no form of the command disables the associated type of traffic sampling on the associated interface.</p> |
| Default | no sampling |
| Parameters | <p>unicast — Specifies that the sampling command will control the sampling of unicast traffic on the associated interface/SAP.</p> <p>multicast — Specifies that the sampling command will control the sampling of multicast traffic on the associated interface/SAP.</p> <p>type — Specifies whether the traffic sampling is based on an acl match, or all traffic entering or exiting the associated interface.</p> <p>Values</p> <p> acl — Specifies that the sampled traffic is controlled via an IP traffic filter entry with the action “filter-sample” configured.</p> <p> interface — Specifies that all traffic entering or exiting the interface is subject to sampling.</p> |

direction — Specifies the direction to collect traffic flow samples.

Values ingress-only — Enables ingress sampling only on the associated interface.
 egress-only — Enables egress sampling only on the associated interface.
 both — Enables both ingress and egress cflowd sampling.

cpu-protection

Syntax **cpu-protection** *policy-id*
 no cpu-protection

Context config>router>if

Description This command assigns an existing CPU protection policy for the interface. The CPU protection policies are configured in the **config>sys>security>cpu-protection>policy** *cpu-protection-policy-id* context.

Default cpu-protection 255

Parameters *policy-id* — Specifies an existing CPU protection policy

Values 1 to 255

dist-cpu-protection

Syntax **dist-cpu-protection** *policy-name*
 no dist-cpu-protection

Context config>router>if

Description This command assigns a Distributed CPU Protection (DCP) policy to the SAP. Only a valid created DCP policy can be assigned to a SAP or a network interface (note that this rule does not apply to templates such as an msap-policy).

Default If no dist-cpu-protection policy is assigned to a router interface, then, the default network DCP policy (_default-network-policy) is used.

 If no DCP functionality is desired on the router interface then an empty DCP policy can be created and explicitly assigned to the router interface.

Parameters *policy-name* — Specifies the name of the DCP policy up to 32 characters in length

enable-ingress-stats

Syntax **[no] enable-ingress-stats**

| | |
|--------------------|---|
| Context | <pre>config>router>if config>service>ies >if config>service>vpn>if config>service>ies>sub-if>grp-if config>service>vpn>sub-if>grp-if</pre> |
| Description | <p>This command enables the collection of ingress interface IP stats. This command is only applicable to IP statistics, and not to uRPF statistics.</p> <p>If enabled, then the following statistics are collected:</p> <ul style="list-style-type: none"> • IPv4 offered packets • IPv4 offered octets • IPv6 offered packets • IPv6 offered octets <p>Octet statistics for IPv4 and IPv6 bytes at IP interfaces include the Layer 2 frame overhead.</p> |
| Default | no enable-ingress-stats |

enable-mac-accounting

| | |
|--------------------|--|
| Syntax | [no] enable-mac-accounting |
| Context | config>router>if |
| Description | This command enables MAC Accounting functionality for the interface. |
| Default | no enable-mac-accounting |

eth-cfm

| | |
|--------------------|---|
| Syntax | eth-cfm |
| Context | config>router>if |
| Description | This command enables the context to configure ETH-CFM parameters. |

mep

| | |
|--------------------|---|
| Syntax | [no] mep mep-id domain md-index association ma-index |
| Context | config>router>if>eth-cfm |
| Description | <p>This command provisions an 802.1ag maintenance endpoint (MEP).</p> <p>The no form of the command deletes the MEP.</p> |

| | |
|-------------------|--|
| Parameters | <i>mep-id</i> — Specifies the MEP identifier. Values 1 to 81921 <i>md-index</i> — Specifies the maintenance domain (MD) index value. Values 1 to 4294967295 <i>ma-index</i> — Specifies the maintenance association (MA) index value. Values 1 to 4294967295 |
|-------------------|--|

collect-lmm-fc-stats

| | |
|--------------------|--|
| Syntax | collect-lmm-fc-stats |
| Context | config>router>if>eth-cfm>mep |
| Description | <p>This command enables the context to configure per-forwarding class (FC) LMM information collection.</p> <p>This command is mutually exclusive with the collect-lmm-stats command when there is entity resource contention.</p> |

fc

| | |
|--------------------|---|
| Syntax | fc <i>fc-name</i> [<i>fc-name</i> ... (up to 8 max)] no fc |
| Context | config>router>if>eth-cfm>mep>collect-lmm-fc-stats |
| Description | <p>This command creates individual counters for the specified FCs without regard for profile. All countable packets that match a configured FC, regardless of profile, will be included in this counter.</p> <p>A differential is performed when this command is re-entered. Omitted FCs will stop counting, newly added FCs will start counting, and unchanged FCs will continue to count.</p> <p>Up to eight FCs may be specified. An FC that is specified as part of this command for this specific context cannot be specified as a profile-aware FC using the fc-in-profile command under the same context.</p> <p>The no form of the command removes all previously defined FCs and stops counting for those FCs.</p> |
| Default | no fc |

| | |
|-------------------|--|
| Parameters | <p><i>fc-name</i> — Specifies the name of the FC for which to create an individual profile-unaware counter. In order for the counter to be used, the config>oam-pm>session>ethernet>priority command must be configured with a numerical value representing the FC name (7 = NC, 6 = H1, 5 = EF, 4 = H2, 3 = L1, 2 = AF, 1 = L2, 0 = BE), and the config>oam-pm>session>ethernet>lmm>enable-fc-collection command must be enabled.</p> <p>Values nc, h1, ef, h2, l1, af, l2, be</p> |
|-------------------|--|

fc-in-profile

| | |
|--------------------|---|
| Syntax | <p>fc-in-profile <i>fc-name</i> [<i>fc-name</i> ... (up to 8 max)]</p> <p>no fc-in-profile</p> |
| Context | config>router>if>eth-cfm>mep>collect-lmm-fc-stats |
| Description | <p>This command creates individual counters for the specified FCs with regard for profile. All countable packets that match a configured FC and are deemed to be in-profile will be included in this counter.</p> <p>A differential is performed when this command is re-entered. Omitted FCs will stop counting, newly added FCs will start counting, and unchanged FCs will continue to count.</p> <p>Up to eight FCs may be specified. An FC that is specified as part of this command for this specific context cannot be specified as a profile-unaware FC using the fc command under the same context.</p> <p>The no form of the command removes all previously defined FCs and stops counting for those FCs.</p> |
| Default | no fc-in-profile |
| Parameters | <p><i>fc-name</i> — Specifies the name of the FC for which to create an individual profile-aware counter. In order for the counter to be used, the config>oam-pm>session>ethernet>priority command must be configured with a numerical value representing the FC name (7 = NC, 6 = H1, 5 = EF, 4 = H2, 3 = L1, 2 = AF, 1 = L2, 0 = BE), and the config>oam-pm>session>ethernet>lmm>enable-fc-collection command must be enabled.</p> <p>Values nc, h1, ef, h2, l1, af, l2, be</p> |

grace

| | |
|--------------------|--|
| Syntax | grace |
| Context | config>router>if>eth-cfm>mep |
| Description | This command enables the context to configure Nokia ETH-CFM Grace and ITU-T Y.1731 ETH-ED expected defect functional parameters. |

eth-ed

| | |
|--------------------|--|
| Syntax | eth-ed |
| Context | config>router>if>eth-cfm>mep>grace |
| Description | This command enables the context to configure ITU-T Y.1731 ETH-ED expected defect functional parameters. |

max-rx-defect-window

| | |
|--------------------|--|
| Syntax | max-rx-defect-window <i>seconds</i> no max-rx-defect-window |
| Context | config>router>if>eth-cfm>mep>grace>eth-ed |
| Description | <p>This command limits the duration of the received ETH-ED expected defect window to the lower value of either the received value from the peer or this parameter.</p> <p>The no form of the command removes the limitation, and any valid defect window value received from a peer MEP in the ETH-ED PDU will be used.</p> |
| Default | no max-rx-defect-window |
| Parameters | <i>seconds</i> — Specifies the duration, in seconds, of the maximum expected defect window. Values 1 to 86400 |

priority

| | |
|--------------------|--|
| Syntax | priority <i>priority</i> no priority |
| Context | config>router>if>eth-cfm>mep>grace>eth-ed |
| Description | <p>This command sets the priority bits and determines the forwarding class based on the mapping of priority to FC.</p> <p>The no form of the command disables the local priority configuration and sets the priority to the ccm-ltm-priority associated with this MEP.</p> |
| Default | no priority |
| Parameters | <i>priority</i> — Specifies the priority bit. Values 0 to 7 |

rx-eth-ed

| | |
|--------------------|--|
| Syntax | [no] rx-eth-ed |
| Context | config>router>if>eth-cfm>mep>grace>eth-ed |
| Description | <p>This command enables the reception and processing of the ITU-T Y.1731 ETH-ED PDU on the MEP.</p> <p>The no form of the command disables the reception of the ITU-T Y.1731 ETH-ED PDU on the MEP.</p> |
| Default | rx-eth-ed |

tx-eth-ed

| | |
|--------------------|---|
| Syntax | [no] tx-eth-ed |
| Context | config>router>if>eth-cfm>mep>grace>eth-ed |
| Description | <p>This command enables the transmission of the ITU-T Y.1731 ETH-ED PDU from the MEP when a system soft reset notification is received for one or more cards.</p> <p>The config>eth-cfm>system>grace-tx-enable command must be configured to instruct the system that the node is capable of transmitting expected defect windows to the peers. Only one form of ETH-CFM grace (Nokia ETH-CFM Grace or ITU-T Y.1731 ETH-ED) may be transmitted.</p> <p>The no form of the command disables the transmission of the ITU-T Y.1731 ETH-ED PDU from the MEP.</p> |
| Default | no tx-eth-ed |

eth-vsm-grace

| | |
|--------------------|--|
| Syntax | eth-vsm-grace |
| Context | config>router>if>eth-cfm>mep>grace |
| Description | This command enables the context to configure Nokia ETH-CFM Grace functional parameters. |

rx-eth-vsm-grace

| | |
|----------------|--|
| Syntax | [no] rx-eth-vsm-grace |
| Context | config>router>if>eth-cfm>mep>grace>eth-vsm-grace |

| | |
|--------------------|---|
| Description | <p>This command enables the reception and processing of the Nokia ETH-CFM Grace PDU on the MEP.</p> <p>The Nokia Grace function is a vendor-specific PDU that informs MEP peers that the local node may be entering a period of expected defect.</p> <p>The no form of the command disables the reception of the Nokia ETH-CFM Grace PDU on the MEP.</p> |
| Default | rx-eth-vsm-grace |

tx-eth-vsm-grace

| | |
|--------------------|--|
| Syntax | [no] tx-eth-vsm-grace |
| Context | config>router>if>eth-cfm>mep>grace>eth-vsm-grace |
| Description | <p>This command enables the transmission of the Nokia ETH-CFM Grace PDU from the MEP when a system soft reset notification is received for one or more cards.</p> <p>The Nokia Grace function is a vendor-specific PDU that informs MEP peers that the local node may be entering a period of expected defect.</p> <p>The config>eth-cfm>system>grace-tx-enable command must be configured to instruct the system that the node is capable of transmitting expected defect windows to the peers. Only one form of ETH-CFM grace (Nokia ETH-CFM Grace or ITU-T Y.1731 ETH-ED) may be transmitted.</p> <p>The no form of the command disables the transmission of the Nokia ETH-CFM Grace PDU from the MEP.</p> |
| Default | tx-eth-vsm-grace |

lbm-svc-act-responder

| | |
|--------------------|--|
| Syntax | [no] lbm-svc-act-responder |
| Context | config>router>if>eth-cfm>mep |
| Description | <p>This command enables the MEP to process service activation streams encapsulated in ETH-CFM LBM frames that are directed to the MEP. The MEP will be allocated additional resources to rapidly respond to a high-speed stream of LBM messages. A MEP created with this option will not validate any TLVs, will not validate the ETH-LBM MAC Address, and will not increment or compute any loopback statistics. Statistical computation and reporting is the responsibility of the test head-end. The ETH-CFM level of the high speed ETH-LBM stream</p> |

must match the level of a MEP configured with this command. It must not target any lower ETH-CFM level the MEP will terminate. When the service activation test is complete, the MEP may be returned to standard processing by removing this command. If there is available bandwidth, the MEP will respond to other ETH-CFM PDUs, such as ETH-DMM marker packets, using standard processing.

The interaction between this command and the **tools perform service id service-id loopback eth** command must be carefully considered. It is recommended that either the **lbm-svc-act-responder** or the **tools perform service id service-id loopback eth** command be used at any given time within a service. If both commands must be configured, and the target reflection point is the MAC Swap Loopback function, the inbound stream of data must not include ETH-CFM traffic that is equal to or lower than the domain level of any configured MEP which would otherwise extract and process the ETH-CFM message. If the reflection target is a MEP configured with the **lbm-svc-act-responder** option, the mode (ingress or egress) of the SAP or SDP specified with this tools command and the MEP **direction** (up or down) must match when the functions are enabled on the same reflection point, and the domain level of the inbound ETH-LBM must be the same as that of the MEP configured with the **lbm-svc-act-responder** option. At no time should the two functions be conflicting with each other along the path of the stream. This conflict would lead to unpredictable and possibly destabilizing situations.

The **no** form of the command reverts to MEP LBM standard processing.

Default no lbm-svc-act-responder

if-attribute

Syntax **if-attribute**

Context config>router>if

Description This command adds and removes interface attributes.

local-proxy-arp

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

Context config>router>if

Description This command enables local proxy ARP on the interface.

Default no local-proxy-arp

ip-mtu

Syntax **ip-mtu** *octets*

| | |
|--------------------|---|
| | no ip-mtu |
| Context | config>router>if |
| Description | <p>This command configures the IP maximum transmit unit (packet) for the associated router IP interface.</p> <p>The configured IP-MTU cannot be larger than the calculated IP MTU based on the port MTU configuration.</p> <p>The MTU that will be used is:</p> <p>MINIMUM((Port_MTU - EtherHeaderSize), (Configured ip-mtu))</p> <p>The no form of the command returns the associated IP interfaces MTU to its default value, which is calculated, based on the port MTU setting. (For Ethernet ports this will typically be 1554.)</p> |
| Default | no ip-mtu |
| Parameters | <p><i>octets</i> — Specifies the IP MTU value associated with the IP interface, specified in octets.</p> <p>Values 512 to 9000</p> |

ip-tunnel

| | |
|--------------------|--|
| Syntax | ip-tunnel |
| Context | config>router>if |
| Description | <p>This command enables the context to configure parameters for an IP tunnel on a control-channel loopback interface. The default encapsulation is IP/GRE. The local end tunnel IP address will be configured using the interface primary IP address.</p> <p>The ip-tunnel command can only be configured on control-channel loopback interfaces.</p> |

remote-ip

| | |
|--------------------|---|
| Syntax | remote-ip <i>ip-address</i> no remote-ip |
| Context | config>router>if>ip-tunnel |
| Description | <p>This command configures the far-end IP address for an IP/GRE tunnel used by a control-channel loopback interface. The address refers to the “to” address of the outer IP header in the encapsulation.</p> |
| Default | no remote-ip |

Parameters *ip-address* — Specifies an IPv4 address.
Values a.b.c.d

lag-link-map-profile

Syntax **lag-link-map-profile** *link-map-profile-id*
 no lag-link-map-profile

Context config>router>if

Description This command assigns a preconfigured lag link map profile to a SAP/network interface configured on a LAG or a PW port that exists on a LAG. Once assigned/unassigned, the SAP/network interface egress traffic will be re-hashed over LAG as required by the new configuration.

 The **no** form of this command reverts the SAP/network interface to use per-flow, service or link hash as configured for the service/LAG.

Default no lag-link-map-profile

Parameters *link-map-profile-id* — An integer from 1 to 32 that defines a unique lag link map profile on which the LAG the SAP/network interface exist.

lag-per-link-hash

Syntax **lag-per-link-hash class** {1 | 2 | 3} **weight** [1..1024]
 no lag-per-link-hash

Context config>router>if

Description This command configures weight and class to this interface to be used on LAG egress when the LAG uses weighted per-link-hash.

 The **no** form of this command restores the default configuration.

Default no lag-per-link-hash (equivalent to weight 1 class 1)

ldp-shortcut

Syntax **[no] ldp-shortcut**

Context config>router

Description This command enables the resolution of IGP routes using LDP LSP across all network interfaces participating in the IS-IS and OSPF routing protocol in the system.

When LDP shortcut is enabled, LDP populates the routing table with next-hop entries corresponding to all prefixes for which it activated an LDP FEC. For a given prefix, two route entries are populated in the system routing table. One route corresponds to the LDP shortcut next-hop and has an owner of LDP. The other route is the regular IP next-hop. The LDP shortcut next-hop always has preference over the regular IP next-hop for forwarding user packets and specified control packets over a given outgoing interface to the route next-hop.

All user and specified control packets for which the longest prefix match in RTM yields the FEC prefix will be forwarded over the LDP LSP.

When an IPv4 packet is received on an ingress network interface, a subscriber IES interface, or a regular IES interface, the lookup of the packet by the ingress IOM, IMMM, or XMA will result in the packet being sent labeled with the label stack corresponding to the NHLFE of the LDP LSP when the preferred RTM entry corresponds to an LDP shortcut.

If the preferred RTM entry corresponds to an IP next-hop, the IPv4 packet is forwarded unlabeled.

When ECMP is enabled and multiple equal-cost next-hops exit for the IGP route, the ingress IOM, IMMM, or XMA will spray the packets for this route based on hashing routine currently supported for IPv4 packets. When the preferred RTM entry corresponds to an LDP shortcut route, spraying will be performed across the multiple next-hops for the LDP FEC. The FEC next-hops can either be direct link LDP neighbors or T-LDP neighbors reachable over RSVP LSPs in the case of LDP-over-RSVP but not both.

When the preferred RTM entry corresponds to a regular IP route, spraying will be performed across regular IP next-hops for the prefix.

The **no** form of this command disables the resolution of IGP routes using LDP shortcuts.

Default no ldp-shortcut

ldp-sync-timer

Syntax ldp-sync-timer *seconds* [**end-of-lib**]
no ldp-sync-timer

Context config>router>if

Description This command enables synchronization of an IGP and LDP. When a link is restored after a failure, the IGP sets the link cost to infinity and advertises it. The actual value advertised in OSPF is 0xFFFF (65535). The actual value advertised in IS-IS regular metric is 0x3F (63) and in IS-IS wide-metric is 0xFFFFFE (16777214). This feature is not supported on RIP interfaces.

If an interface belongs to both IS-IS and OSPF, a physical failure will cause both IGPs to advertise an infinite metric and to follow the IGP-LDP synchronization procedures. If only one IGP bounces on this interface or on the system, then only the affected IGP advertises the infinite metric and follows the IGP-LDP synchronization procedures.

Next, an LDP Hello adjacency is brought up with the neighbor. The LDP synchronization timer is started by the IGP when the LDP session to the neighbor is up over the interface. This is to allow time for the label-FEC bindings to be exchanged.

When the LDP synchronization timer expires, the link cost is restored and is readvertised. The IGP will announce a new best next hop and LDP will use it if the label binding for the neighbor's FEC is available.

If the user changes the cost of an interface, the new value is advertised at the next flooding of link attributes by the IGP. However, if the LDP synchronization timer is still running, the new cost value will only be advertised after the timer expires. The new cost value will also be advertised after the user executes any of the following commands:

- **tools>perform>router>isis>ldp-sync-exit**
- **tools>perform>router>ospf>ldp-sync-exit**
- **config>router>if>no ldp-sync-timer**
- **config>router>ospf>disable-ldp-sync**
- **router>isis>disable-ldp-sync**

If the user changes the value of the LDP synchronization timer parameter, the new value will take effect at the next synchronization event. If the timer is still running, it will continue to use the previous value.

If parallel links exist to the same neighbor, then the bindings and services should remain up as long as there is one interface that is up. However, the user-configured LDP synchronization timer still applies on the interface that failed and was restored. In this case, the router will only consider this interface for forwarding after the IGP readvertises its actual cost value.

The LDP Sync Timer State is not always synchronized across to the standby CPM. Therefore, after an activity switch, the timer state might not be same as it was on the previously active CPM.

If the **end-of-lib** option is configured, then the system will start the LDP synchronization timer as usual. If the LDP End of LIB Typed Wildcard FEC messages are received for every FEC type negotiated for a given session to an LDP peer for that IGP interface, the **ldp-sync-timer** is terminated early and the IGP link cost is restored. If the **ldp-sync-timer** expires before the LDP End of LIB messages are received for every negotiated FEC type, then the system will restore the IGP link cost. The **end-of-lib** option is disabled by default.

The **no** form of this command disables IGP-LDP synchronization and deletes the configuration.

| | |
|-------------------|---|
| Default | no ldp-sync-timer |
| Parameters | <i>seconds</i> — Specifies the time interval for the IGP-LDP synchronization timer. |
| Values | 1 to 1800 |

end-of-lib — Specifies that the system should terminate the **ldp-sync-timer** early if the LDP End of LIB Typed Wildcard FEC messages are received for every FEC type negotiated for a given session to an LDP peer for that IGP interface.

load-balancing

| | |
|--------------------|---|
| Syntax | load-balancing |
| Context | config>router>if |
| Description | This command enables the load-balancing context to configure interface per-flow load balancing options that will apply to traffic entering this interface and egressing over a LAG/ECMP on system-egress. This is a per interface setting. For load-balancing options that can also be enabled on the system level, the options enabled on the interface level overwrite system level configurations. |

egr-ip-load-balancing

| | |
|--------------------|---|
| Syntax | egr-ip-load-balancing {source destination inner-ip} no egr-ip-load-balancing |
| Context | config>router>if>load-balancing |
| Description | <p>This command specifies whether to include source address or destination address or both in LAG/ECMP hash on IP interfaces. Additionally, when l4-load-balancing is enabled the command applies also to inclusion of source/destination port in the hash inputs.</p> <p>The no form of this command includes both source and destination parameters.</p> |
| Default | no egr-ip-load-balancing |
| Parameters | <p>source — Specifies using source address and (if l4-load balancing is enabled) source port in the hash, ignore destination address/port</p> <p>destination — Specifies using destination address and (if l4-load balancing is enabled) destination port in the hash, ignore source address/port.</p> <p>inner-ip — Specifies use of the inner IP header parameters instead of outer IP header parameters in LAG/ECMP hash for IPv4 encapsulated traffic.</p> |

lsr-load-balancing

| | |
|----------------|--|
| Syntax | lsr-load-balancing <i>hashing-algorithm</i> no lsr-load-balancing |
| Context | config>router>if>load-balancing |

| | |
|--------------------|---|
| Description | This command specifies whether the IP header is used in the LAG and ECMP LSR hashing algorithm. This is the per interface setting. |
| Default | no lsr-load-balancing |
| Parameters | <p>lbl-only — Specifies that only the label is used in the hashing algorithm</p> <p>lbl-ip — Specifies that only the IP header is included in the hashing algorithm.</p> <p>ip-only — Specifies that only the IP header is used exclusively in the hashing algorithm</p> <p>eth-encap-ip — Specifies that the hash algorithm parses down the label stack (up to 3 labels supported) and once it hits the bottom, the stack assumes Ethernet II non-tagged header follows. At the expected Ethertype offset location, algorithm checks whether the value present is IPv4/v6 (0x0800 or 0x86DD). If the check passes, the hash algorithm checks the first nibble at the expected IP header location for IPv4/IPv6 (0x0100/0x0110). If the secondary check passes, the hash is performed using IP SA/DA fields in the expected IP header; otherwise (any of the check failed) label-stack hash is performed.</p> <p>lbl-ip-l4-teid — Specifies that this hashing algorithm hashes based on label, IP header, Layer 4 header and GTP header (TEID) in order. The algorithm uses all the supported headers that are found in the header fragment of incoming traffic.</p> |

spi-load-balancing

| | |
|--------------------|---|
| Syntax | [no] spi-load-balancing |
| Context | config>router>if>load-balancing |
| Description | <p>This command enables use of the SPI in hashing for ESP/AH encrypted IPv4/v6 traffic. This is a per interface setting.</p> <p>The no form disables the SPI function.</p> |
| Default | no spi-load-balancing |

teid-load-balancing

| | |
|--------------------|---|
| Syntax | [no] teid-load-balancing |
| Context | config>router>if>load-balancing |
| Description | This command enables inclusion of TEID in hashing for GTP-U/C encapsulates traffic for GTPv1/GTPv2. The no form of this command ignores TEID in hashing. |
| Default | no teid-load-balancing |

loopback

| | |
|--------------------|---|
| Syntax | [no] loopback |
| Context | config>router>if |
| Description | This command configures the interface as a loopback interface. The vas-if-type and loopback commands are mutually exclusive |
| Default | Not enabled |

mac

| | |
|--------------------|---|
| Syntax | mac <i>ieee-mac-addr</i> no mac |
| Context | config>router>if |
| Description | <p>This command assigns a specific MAC address to an IP interface. Only one MAC address can be assigned to an IP interface. When multiple mac commands are entered, the last command overwrites the previous command.</p> <p>The no form of the command returns the MAC address of the IP interface to the default value.</p> |
| Default | IP interface has a system-assigned MAC address. |
| Parameters | <i>ieee-mac-addr</i> — Specifies the 48-bit MAC address for the IP interface in the form <i>aa:bb:cc:dd:ee:ff</i> or <i>aa-bb-cc-dd-ee-ff</i> , where <i>aa</i> , <i>bb</i> , <i>cc</i> , <i>dd</i> , <i>ee</i> and <i>ff</i> are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses. |

network-domain

| | |
|--------------------|---|
| Syntax | network-domain <i>network-domain-name</i> no network-domain |
| Context | config>router>if |
| Description | <p>This command assigns a given interface to a given network-domain. The network-domain is then taken into account during sap-ingress queue allocation for VPLS SAP.</p> <p>The network-domain association can only be done in a base-routing context. Associating a network domain with an loop-back or system interface will be rejected. Associating a network-domain with an interface that has no physical port specified will be accepted, but will have no effect as long as a corresponding port, or LAG, is defined.</p> <p>Single interfaces can be associated with multiple network-domains.</p> |

Default per default “default” network domain is assigned

ntp-broadcast

Syntax **[no] ntp-broadcast**

Context config>router>if

Description This command enables SNTP broadcasts received on the IP interface. This parameter is only valid when the SNTP **broadcast-client** global parameter is configured.

The **no** form of the command disables SNTP broadcast received on the IP interface.

Default no ntp-broadcast

port

Syntax **port** *port-name*
no port

Context config>router>if

Description This command creates an association with a logical IP interface and a physical port.

An interface can also be associated with the system (loopback address).

The command returns an error if the interface is already associated with another port or the system. In this case, the association must be deleted before the command is re-attempted. The *port-id* or *port-id* for Ethernet ports can be in one of the following forms:

Ethernet interfaces

If the card in the slot has MDAs/XMAs, *port-id* is in the *slot_number/MDA* or *XMA_number/port_number* format; for example, **1/1/3** specifies port 3 of the MDA/XMA installed in MDA/XMA slot 1 on the card installed in chassis slot 1.

SONET/SDH interfaces

When the *port-id* represents a POS interface, the *port-id* must include the *channel-id*. The POS interface must be configured as a **network** port.

The **no** form of the command deletes the association with the port. The **no** form of this command can only be performed when the interface is administratively down.

Default No port is associated with the IP interface.

Parameters *port-name* — The physical port identifier to associate with the IP interface.

Values The following values apply to the 7750 SR:

| | | | |
|------------------|---|---------------------------------|-----------------------|
| <i>port-name</i> | <i>port-id[:encap-val]</i> | | |
| | encap-val | 0 | for null |
| | | 0..4094 | for dot1q |
| | | 0..4094.* | for qinq |
| <i>port-id</i> | <i>slot/mda/port[.channel]</i> | | |
| | eth-sat-id | <i>esat-id/slot/port</i> | |
| | | esat | keyword |
| | | <i>id</i> | 1 to 20 |
| | pxc-id | <i>pxc-id.sub-port</i> | |
| | | pxc | keyword |
| | | <i>id</i> | 1 to 64 |
| | | <i>sub-port</i> | a, b |
| | <i>bundle-id</i> - <i>bundle-type-slot/mda.bundle-num</i> | | |
| | | bundle | keyword |
| | | type | ima, fr, ppp |
| | | bundle-num | 1..336 |
| | bpgrp-id | <i>bpgrp-type-bpgrp-num</i> | |
| | | bpgrp | keyword |
| | | type | ima, ppp |
| | | bpgrp-num | 1 to 2000 |
| | aps-id | <i>aps-group-id[.channel]</i> | |
| | | aps | keyword |
| | | group-id | 1 to 64 |
| | ccag-id | <i>ccag-id.path-id[cc-type]</i> | |
| | | ccag | keyword |
| | | id | 1..8 |
| | | path-id | a, b |
| | | cc-type | .sap-net, .net-sap |
| | lag-id | <i>lag-id</i> | |
| | | lag | keyword |
| | | id | 1 to 800 |
| <i>port-id</i> | <i>slot/mda/ port[.channel]</i> | | |
| | eth-sat-id | <i>esat-id/slot/port</i> | |
| | | esat | keyword |
| | | <i>id</i> | 1 to 20 |
| | pxc-id | <i>pxc-id.sub-port</i> | |

| | | | |
|---------|-------------------------|---------------------------------|--------------------|
| | | pxc | keyword |
| | | <i>id</i> | 1 to 64 |
| | | <i>sub-port</i> | a, b |
| | bundle-id | bundle-type-slot/mda.bundle-num | |
| | | bundle | keyword |
| | | type | ima, ppp |
| | | bundle-num | 1 to 336 |
| | bpgrp-id | bpgrp-type-bpgrp-num | |
| | | bpgrp | keyword |
| | | type | ima, ppp |
| | | bpgrp-num | 1 to 256 |
| | aps-id | aps-group-id[.channel] | |
| | | aps | keyword |
| | | group-id | 1 to 16 |
| | lag-id | lag-id | |
| | | lag | keyword |
| | | id | 1 to 64 |
| port-id | slot/mda/port[.channel] | | |
| | eth-sat-id | <i>esat-id/slot/port</i> | |
| | | esat | keyword |
| | | <i>id</i> | 1 to 20 |
| | pxc-id | <i>pxc-id.sub-port</i> | |
| | | pxc | keyword |
| | | <i>id</i> | 1 to 64 |
| | | <i>sub-port</i> | a, b |
| | ccag-id | ccag-id.path-id[cc-type] | |
| | | ccag | keyword |
| | | id | 1 to 8 |
| | | path-id | a, b |
| | | cc-type | .sap-net, .net-sap |
| | lag-id | lag-id | |
| | | lag | keyword |
| | | id | 1 to 200 |
| | gtg-id | gmpls-tun-grp-id | |
| | | gmpls-tun-grp | keyword |
| | | id | 1 to 1024 |

Values The following values apply to the 7450 ESS:

| | | | |
|------------|---------------------------------|---------|--------------------|
| port-id | <i>slot/mda/port[.channel]</i> | | |
| eth-sat-id | <i>esat-id/slot/port</i> | | |
| | <i>esat</i> | keyword | |
| | <i>id</i> | | 1 to 20 |
| pxc-id | <i>pxc-id.sub-port</i> | | |
| | <i>pxc</i> | keyword | |
| | <i>id</i> | | 1 to 64 |
| | <i>sub-port</i> | | a, b |
| ccag-id | <i>ccag-id.path-id[cc-type]</i> | | |
| | <i>ccag</i> | keyword | |
| | <i>id</i> | | 1 to 8 |
| | <i>path-id</i> | | a, b |
| | <i>cc-type</i> | | .sap-net, .net-sap |
| lag-id | <i>lag-id</i> | | |
| | <i>lag</i> | keyword | |
| | <i>id</i> | | 1 to 800 |
| gtg-id | <i>gmpls-tun-grp-id</i> | | |
| | <i>gmpls-tun-grp</i> | keyword | |
| | <i>id</i> | | 1 to 200 |

proxy-arp-policy

| | |
|--------------------|---|
| Syntax | [no] proxy-arp-policy <i>policy-name</i> [<i>policy-name...</i> (up to 5 max)] |
| Context | config>router>if |
| Description | <p>This command enables and configure proxy ARP on the interface and specifies an existing policy-statement to analyze match and action criteria that controls the flow of routing information to and from a given protocol, set of protocols, or a specific neighbor. The policy-name is configured in the config>router>policy-options context.</p> <p>Use proxy ARP so the router responds to ARP requests on behalf of another device. Static ARP is used when a router needs to know about a device on an interface that cannot or does not respond to ARP requests. Therefore, the router configuration can state that if it has a packet that has a certain IP address to send it to the corresponding ARP address.</p> |
| Default | no proxy-arp-policy |
| Parameters | <p><i>policy-name</i> — The export route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. The specified policy names must already be defined.</p> |

ptp-hw-assist

| | |
|--------------------|---|
| Syntax | [no] ptp-hw-assist |
| Context | config>router>if |
| Description | <p>This command configures the 1588 port based timestamping assist function for the interface. Various checks are performed to ensure that this feature can be enabled. If a check fails:</p> <ul style="list-style-type: none"> • The command is blocked/rejected with an appropriate error message. • If the SAP configuration of the interface is removed, the ptp-hw-assist configuration will be removed. • If the IPv4 address configuration of the interface is removed, the ptp-hw-assist configuration will be removed. <p>The port will validate the destination IP address on received 1588 messages. If the 1588 messages are sent to a loopback address within the node rather than the address of the interface, then the loopback address must be configured in the configure>system>security>source-address application ptp context.</p> |
| Default | no ptp-hw-assist |

qos-route-lookup

| | |
|--------------------|--|
| Syntax | qos-route-lookup [source destination] no qos-route-lookup |
| Context | config>router>if config>router>if>ipv6 |
| Description | <p>This command enables QoS classification of the ingress IP packets on an interface based on the QoS information associated with routes in the forwarding table.</p> <p>If the optional destination parameter is specified and the destination address of an incoming IP packet matches a route with QoS information the packet is classified to the fc and priority associated with that route, overriding the fc and priority/profile determined from the sap-ingress or network qos policy associated with the IP interface. If the destination address of the incoming packet matches a route with no QoS information the fc and priority of the packet remain as determined by the sap-ingress or network qos policy.</p> <p>If the optional source parameter is specified and the source address of an incoming IP packet matches a route with QoS information the packet is classified to the fc and priority associated with that route, overriding the fc and priority/profile determined from the sap-ingress or network qos policy associated with the IP interface. If the source address of the incoming packet matches a route with no QoS information the fc and priority of the packet remain as determined by the sap-ingress or network qos policy.</p> <p>If neither the optional source or destination parameter is present, then the default is destination address matching.</p> |

The functionality enabled by the `qos-route-lookup` command can be applied to IPv4 packets or IPv6 packets on an interface, depending on whether it is present at the interface context (applies to IPv4) or the `interface>ipv6` context (applies to IPv6). Subscriber management group interfaces for the 7750 SR and 7450 ESS also do not support the source QPPB option.

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

Default no qos-route-lookup

Parameters **source** — Enables QoS classification of incoming IP packets based on the source address matching a route with QoS information.

destination — Enables QoS classification of incoming IP packets based on the destination address matching a route with QoS information.

qos

Syntax **qos** *network-policy-id* [**egress-port-redirect-group** *queue-group-name*] [**egress-instance** *instance-id*] [**ingress-fp-redirect-group** *queue-group-name* **ingress-instance** *instance-id*]

no qos

Context config>router>if

Description This command associates a network Quality of Service (QoS) policy with a network IP interface. Only one network QoS policy can be associated with an IP interface at one time. Attempts to associate a second QoS policy return an error.

Associating a network QoS policy with a network interface is useful for the following purposes:

- To apply classification rules for determining the forwarding-class and profile of ingress packets on the interface.
- To associate ingress packets on the interface with a queue-group instance applied to the ingress context of the interface's forwarding plane (FP). (This is only applicable to interfaces on IOM3 and later cards.) The referenced ingress queue-group instance may have policers defined in order to rate limit ingress traffic on a per-forwarding class (and forwarding type: unicast vs. multicast) basis.
- To perform 802.1p, DSCP, IP precedence and/or MPLS EXP re-marking of egress packets on the interface.
- To associate egress packets on the interface with a queue-group instance applied to the egress context of the interface's port. The referenced egress queue-group instance may have policers and/or queues defined in order to rate limit egress traffic on a per-forwarding class basis.

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

Default no qos

| | |
|-----------------------------------|--|
| Parameters | <i>network-policy-id</i> — An existing network policy ID to associate with the IP interface. |
| Values | 1 to 65535 |
| egress-port-redirect-group | <i>queue-group-name</i> — This optional parameter specifies the egress queue-group used for all egress forwarding-class redirections specified within the network QoS policy ID. The specified <i>queue-group-name</i> must exist as an egress queue group applied to the egress context of the port associated with the IP interface. |
| egress-instance | <i>instance-id</i> — Since multiple instances of the same egress queue-group can be applied to the same port this optional parameter is used to specify which instance to associate with this specific network IP interface. |
| Values | 1 to 16384 |
| ingress-fp-redirect-group | <i>queue-group-name</i> — This optional parameter specifies the ingress queue-group used for all ingress forwarding-class redirections specified within the network QoS policy ID. The specified <i>queue-group-name</i> must exist as an ingress queue group applied to the ingress context of the forwarding plane associated with the IP interface. |
| ingress-instance | <i>instance-id</i> — Since multiple instances of the same ingress queue-group can be applied to the same forwarding plane this parameter is required to specify which instance to associate with this specific network IP interface. |
| Values | 1 to 16384 |

remote-proxy-arp

| | |
|--------------------|---|
| Syntax | [no] remote-proxy-arp |
| Context | config>router>if |
| Description | This command enables remote proxy ARP on the interface. |
| Default | no remote-proxy-arp |

secondary

| | |
|--------------------|---|
| Syntax | secondary {[<i>ip-address/mask</i> <i>ip-address netmask</i>]} [broadcast { all-ones host-ones }] [igp-inhibit] no secondary <i>ip-addr</i> |
| Context | config>router>if |
| Description | This command assigns additional IP addresses to the interface. Up to 16 total primary and secondary IPv4 and IPv6 addresses can be assigned to network interfaces, and up to 256 to access interfaces. Each address can be configured in an IP address, IP subnet, or broadcast address format. |



Caution: Configurations must not exceed 16 secondary IP addresses when IPsec, GRE, L2TPv3, or IP in IP protocols are active on an access interface.

Parameters

ip-address — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

Values 1.0.0.0 to 223.255.255.255

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

mask-length — The subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-address* from the *mask-length* parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address. Allowed values are integers in the range 1 to 32. A mask length of 32 is reserved for system IP addresses.

Values 1 to 32

mask — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-addr* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that will be used in a logical ‘AND’ function to derive the local subnet of the IP address. A mask of 255.255.255.255 is reserved for system IP addresses.

Values 128.0.0.0 to 255.255.255.255

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

The **all-ones** keyword following the **broadcast** parameter specifies that the broadcast address used by the IP interface for this IP address will be 255.255.255.255, also known as the local broadcast.

The **host-ones** keyword following the **broadcast** parameter specifies that the broadcast address used by the IP interface for this IP address will be the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the *ip-addr* and the *mask-length* or *mask* with all the host bits set to binary 1. This is the default broadcast address used by an IP interface.

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

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

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

igp-inhibit — The secondary IP address should not be recognized as a local interface by the running IGP.

static-arp

| | |
|--------------------|--|
| Syntax | static-arp <i>ip-addr ieee-mac-addr unnumbered</i> no static-arp <i>unnumbered</i> |
| Context | config>router>if |
| Description | <p>This command configures a static Address Resolution Protocol (ARP) entry associating an IP address with a MAC address for the core router instance. This static ARP appears in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface.</p> <p>If an entry for a specific IP address already exists and a new MAC address is configured for the IP address, the existing MAC address is replaced by the new MAC address.</p> <p>The number of static-arp entries that can be configured on a single node is limited to 1000.</p> <p>Static ARP is used when a router needs to know about a device on an interface that cannot or does not respond to ARP requests. Therefore, the router configuration can state that if it has a packet that has a certain IP address to send it to the corresponding ARP address. Use proxy ARP so the router responds to ARP requests on behalf of another device.</p> <p>The no form of the command removes a static ARP entry.</p> |
| Default | No static ARPs are defined. |
| Parameters | <p><i>unnumbered</i> — Specifies the static ARP MAC for an unnumbered interface. Unnumbered interfaces support dynamic ARP. Once this command is configured, it overrides any dynamic ARP.</p> <p><i>ieee-mac-addr</i> — Specifies the 48-bit MAC address for the static ARP in the form <i>aa:bb:cc:dd:ee:ff</i> or <i>aa-bb-cc-dd-ee-ff</i>, where <i>aa</i>, <i>bb</i>, <i>cc</i>, <i>dd</i>, <i>ee</i> and <i>ff</i> are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.</p> |

strip-label

| | |
|--------------------|--|
| Syntax | [no] strip-label |
| Context | config>router>if |
| Description | <p>This command forces packets to be stripped of all (max 5) MPLS labels before the packets are handed over for possible filter (PBR) processing.</p> <p>If the packets do not have an IP header immediately following the MPLS label stack after the strip, they are discarded. Only MPLS encapsulated IP, IGP shortcuts and VPRN over MPLS packets will be processed. However, IPv4 and IPv6 packets that arrive without any labels are supported on an interface with strip-label enabled.</p> <p>This command is supported on:</p> <ul style="list-style-type: none">• Optical ports for the 7750 SR and 7450 ESS• IOM3-XP cards for the 7750 SR and 7450 ESS• Null/Dot1q encaps• Network ports• IPv4• IPv6 <p>The no form of the command removes the strip-label command.</p> <p>In order to associate an interface that is configured with the strip-label parameter with a port, the port must be configured as single-fiber for the command to be valid.</p> <p>Packets that are subject to the strip-label action and are mirrored (using mirrors or lawful interception) will contain the original MPLS labels (and other L2 encapsulation) in the mirrored copy of the packet, as they appeared on the wire, when the mirror-dest type is the default type "ether". If the mirror-dest type is "ip-only", then the mirrored copy of the packet will not contain the original L2 encapsulation or the stripped MPLS labels.</p> |
| Default | no strip-label |

tos-marking-state

| | |
|--------------------|--|
| Syntax | tos-marking-state {trusted untrusted} no tos-marking-state |
| Context | config>router>if |
| Description | <p>This command is used on a network IP interface to alter the default trusted state to a non-trusted state. When unset or reverted to the trusted default, the ToS field will not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set, in which case the egress network interface treats all IES and network IP interface as untrusted.</p> |

When the ingress network IP interface is set to untrusted, all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface. The egress network remarking rules also apply to the ToS field of IP packets routed using IGP shortcuts (tunneled to a remote next-hop). However, the tunnel QoS markings are always derived from the egress network QoS definitions. Egress marking and remarking is based on the internal forwarding class and profile state of the packet once it reaches the egress interface. The forwarding class is derived from ingress classification functions. The profile of a packet is either derived from ingress classification or ingress policing.

The default marking state for network IP interfaces is trusted. This is equivalent to declaring no tos-marking-state on the network IP interface. When undefined or set to tos-marking-state trusted, the trusted state of the interface will not be displayed when using show config or show info unless the detail parameter is given. The **save config** command will not store the default tos-marking-state trusted state for network IP interfaces unless the detail parameter is also specified.

The **no** form of the command is used to restore the trusted state to a network IP interface. This is equivalent to executing the tos-marking-state trusted command.

Default tos-marking-state trusted

Parameters **trusted** — The default prevents the ToS field to not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set

untrusted — Specifies that all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface.

unnumbered

Syntax **unnumbered** [*ip-address* | *ip-int-name*]
no unnumbered

Context config>router>if

Description This command sets an IP interface as an unnumbered interface and specifies the IP address to be used for the interface.

To conserve IP addresses, unnumbered interfaces can be configured. The address used when generating packets on this interface is the *ip-addr* parameter configured.

An error message will be generated if an **unnumbered** interface is configured, and an IP address already exists on this interface.

The **no** form of the command removes the IP address from the interface, effectively removing the unnumbered property. The interface must be **shutdown** before **no unnumbered** is issued to delete the IP address from the interface, or an error message will be generated.

Default no unnumbered

| | |
|-------------------|---|
| Parameters | <i>ip-addr</i> <i>ip-int-name</i> — Optional. The IP address or IP interface name to associate with the unnumbered IP interface in dotted decimal notation. The configured IP address must exist on this node. It is recommended to use the system IP address as it is not associated with a specific interface and is therefore always reachable. The system IP address is the default if no <i>ip-addr</i> or <i>ip-int-name</i> is configured. |
|-------------------|---|

qos-route-lookup

| | |
|--------------------|---|
| Syntax | qos-route-lookup [source destination] no qos-route-lookup |
| Context | config>router>if config>router>if>ipv6 |
| Description | <p>This command enables QoS classification of the ingress IP packets on an interface based on the QoS information associated with routes in the forwarding table.</p> <p>If the optional destination parameter is specified and the destination address of an incoming IP packet matches a route with QoS information the packet is classified to the fc and priority associated with that route, overriding the fc and priority/profile determined from the sap-ingress or network qos policy associated with the IP interface. If the destination address of the incoming packet matches a route with no QoS information the fc and priority of the packet remain as determined by the sap-ingress or network qos policy.</p> <p>If the optional source parameter is specified and the source address of an incoming IP packet matches a route with QoS information the packet is classified to the fc and priority associated with that route, overriding the fc and priority/profile determined from the sap-ingress or network qos policy associated with the IP interface. If the source address of the incoming packet matches a route with no QoS information the fc and priority of the packet remain as determined by the sap-ingress or network qos policy.</p> <p>If neither the optional source or destination parameter is present, then the default is destination address matching.</p> <p>The functionality enabled by the qos-route-lookup command can be applied to IPv4 packets or IPv6 packets on an interface, depending on whether it is present at the interface context (applies to IPv4) or the interface>ipv6 context (applies to IPv6). The ability to specify source address based QoS lookup is not supported for IPv6. Subscriber management group interfaces also do not support the source QPPB option.</p> <p>The no form of the command reverts to the default.</p> |
| Default | destination |
| Parameters | <p>source — Enables QoS classification of incoming IP packets based on the source address matching a route with QoS information.</p> <p>destination — Enables QoS classification of incoming IP packets based on the destination address matching a route with QoS information.</p> |

secure-nd

| | |
|--------------------|---|
| Syntax | [no] secure-nd |
| Context | config>router>if>ipv6 |
| Description | This command enables Secure Neighbor Discovery (SeND) on the IPv6 interface. The no form of the command reverts to the default and disabled SeND. |

allow-unsecured-msgs

| | |
|--------------------|---|
| Syntax | [no] allow-unsecured-msgs |
| Context | config>router>if>ipv6>secure-nd |
| Description | This command specifies whether unsecured messages are accepted. When Secure Neighbor Discovery (SeND) is enabled, only secure messages are accepted by default. The no form of the command disables accepting unsecured messages. |

link-local-modifier

| | |
|--------------------|--|
| Syntax | link-local-modifier <i>modifier</i> [no] link-local-modifier |
| Context | config>router>if>ipv6>secure-nd |
| Description | This command configures the Cryptographically Generated Address (CGA) modifier for link-local addresses. |
| Parameters | <i>modifier</i> — Specifies the modifier in 32 hexadecimal nibbles. Values 0x0 to 0xFFFFFFFF |

public-key-min-bits

| | |
|--------------------|--|
| Syntax | public-key-min-bits <i>bits</i> [no] public-key-min-bits |
| Context | config>router>if>ipv6>secure-nd |
| Description | This command configures the minimum acceptable key length for public keys used in the generation of a Cryptographically Generated Address (CGA). |
| Parameters | <i>bits</i> — Specifies the number of bits. Values 512 to 1024 |

security-parameter

| | |
|--------------------|---|
| Syntax | security-parameter <i>sec</i> [no] security-parameter |
| Context | config>router>if>ipv6>secure-nd |
| Description | This command configures the security parameter used in the generation of a Cryptographically Generated Address (CGA). |
| Parameters | <i>sec</i> — Specifies the security parameter. |
| Values | 0 to 1 |

shutdown

| | |
|--------------------|---|
| Syntax | [no] shutdown |
| Context | config>router>if>ipv6>secure-nd |
| Description | This command enables or disables Secure Neighbor Discovery (SeND) on the interface. |

stale-time

| | |
|--------------------|---|
| Syntax | stale-time <i>seconds</i> no stale-time |
| Context | config>router>ipv6 config>router>if>ipv6 |
| Description | This command configures the time a neighbor discovery cache entry can remain stale before being removed. The no form of the command removes the stale-time value. |
| Default | no stale-time |
| Parameters | <i>seconds</i> — The allowed stale time (in seconds) before a neighbor discovery cache entry is removed. |
| Values | 60 to 65535 |

tcp-mss

| | |
|---------------|--|
| Syntax | tcp-mss <i>mss-value</i> no tcp-mss |
|---------------|--|

| | |
|--------------------|--|
| Context | config>router>if config>router>if>ipv6 |
| Description | <p>This command statically sets the TCP maximum segment size (MSS) for TCP connections originated from the associated IP interface to the specified value.</p> <p>The no form of the command removes the static value and allows the TCP MSS value to be calculated based on the IP MTU value by subtracting the base IP and TCP header lengths from the IP MTU value ($\text{tcp_mss} = \text{ip_mtu} - 40$).</p> |
| Default | no tcp-mss |
| Parameters | <p><i>mss-value</i> — The TCP MSS value that should be used in the TCP SYN packet during the three-way handshake negotiation of a TCP connection.</p> <p>9158 = max-IP_MTU (9198)-40</p> <p>Values 536 to 9158 (IPv4) 1220 to 9138 (IPv6)</p> |

urpf-check

| | |
|--------------------|---|
| Syntax | [no] urpf-check |
| Context | config>router>if config>router>if>ipv6 |
| Description | <p>This command enables unicast RPF (uRPF) Check on this interface.</p> <p>The no form of the command disables unicast RPF (uRPF) Check on this interface.</p> |

urpf-selected-vprns

| | |
|--------------------|---|
| Syntax | urpf-selected-vprns no urpf-selected-vprns |
| Context | config>router>if |
| Description | <p>This command enables uRPF checking of incoming traffic on the network interface for the following packets.</p> <ul style="list-style-type: none"> • Packets associated with the global routing table (base router) context. • Packets associated with VPRNs that have enabled the uRPF check using the config>service>vprn>network>ingress>urpf-check command. <p>If the command is not configured, the default action is to perform uRPF checks for all ingress traffic on the network interface (associated with the base router and all VPRNs) based on the IPv4 and IPv6 urpf-check configuration options of the network interface.</p> |
| Default | no urpf-selected-vprns |

vas-if-type

| | |
|--------------------|---|
| Syntax | vas-if-type { to-from-access to-from-network to-from-both } no vas-if-type |
| Context | config>router>if |
| Description | <p>This command configures the type of a Value Added Service (VAS) facing interface. To change the vas-if-type, the shutdown command is required. The vas-if-type and loopback commands are mutually exclusive.</p> <p>The no form of the command removes the VAS interface type configuration.</p> |
| Default | no vas-if-type |
| Parameters | <p>to-from-access — Used when two separate (to-from-access and to-from-network) interfaces are used for VAS connectivity. For service chaining, traffic arriving from access interfaces (upstream) is redirected to a PBR target reachable over this interface for upstream VAS processing. Downstream traffic after VAS processing must arrive on this interface, so that the traffic is subject to regular routing but is not subject to AA divert, nor egress subscriber PBR.</p> <p>to-from-network — Used when two separate (to-from-access and to-from-network) interfaces are used for VAS connectivity. For service chaining, traffic arriving from network interfaces (downstream) is redirected to a PBR target reachable over this interface for downstream VAS processing. Upstream traffic after VAS processing must arrive on this interface, so that regular routing can be applied.</p> <p>to-from-both — Used when a single interface is used for VAS connectivity (no local-to-local traffic). For service chaining, both traffic arriving from access interfaces and from network interfaces is redirected to a PBR target reachable over this interface for upstream/downstream VAS processing. Traffic after VAS processing must arrive on this interface, so that the traffic is subject to regular routing but is not subject to AA divert, nor to egress subscriber PBR.</p> |

ignore-default

| | |
|--------------------|--|
| Syntax | ignore-default no ignore-default |
| Context | config>router>if>urpf-check config>router>if>ipv6>urpf-check |
| Description | This command configures the uRPF check (if enabled) to ignore default routes for purposes of determining the validity of incoming packets. By default, default routes are considered eligible. |

mode

| | |
|--------------------|---|
| Syntax | mode { strict loose strict-no-ecmp } no mode |
| Context | config>router>if>urpf-check config>router>if>ipv6>urpf-check |
| Description | This command specifies the mode of unicast RPF check. The no form of the command reverts to the default (strict) mode. |
| Default | mode strict |
| Parameters | strict — When specified, uRPF checks whether incoming packet has a source address that matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix. loose — In loose mode, uRPF checks whether incoming packet has source address with a corresponding prefix in the routing table. However, the loose mode does not check whether the interface expects to receive a packet with a specific source address prefix. This object is valid only when urpf-check is enabled. strict-no-ecmp — When a packet is received on an interface in this mode and the SA matches an ECMP route the packet is dropped by uRPF. |

if-attribute

| | |
|--------------------|---|
| Syntax | if-attribute |
| Context | config>router config>router>if config>service>ies>if config>service>vpn>if |
| Description | This command creates the context to configure or apply IP interface attributes such as administrative group (admin-group) or Shared Risk Loss Group (SRLG). |

admin-group

| | |
|--------------------|---|
| Syntax | admin-group <i>group-name</i> value <i>group-value</i> no admin-group <i>group-name</i> |
| Context | config>router>if-attribute |
| Description | This command defines an administrative group (admin-group) that can be associated with an IP or MPLS interface. |

Admin groups, also known as affinity, are used to tag IP and MPLS interfaces that share a specific characteristic with the same identifier. For example, an admin group identifier can represent all links that connect to core routers, or all links that have a bandwidth higher than 10G, or all links that are dedicated to a specific service.

The user first configures locally on each router the name and identifier of each admin group. A maximum of 32 admin groups can be configured per system.

The user then configures the admin group membership of an interface. The user can apply admin groups to a IES, VPRN, network IP, or MPLS interface.

When applied to MPLS interfaces, the interfaces can be included or excluded in the LSP path definition by inferring the admin-group name. CSPF will compute a path that satisfies the admin-group include and exclude constraints.

When applied to IES, VPRN, or network IP interfaces, the interfaces can be included or excluded in the route next-hop selection by inferring the admin-group name in a route next-hop policy template applied to an interface or a set of prefixes.

The following provisioning rules are applied to admin group configuration. The system will reject the creation of an admin-group if it re-uses the same name but with a different group value than an existing group. The system will also reject the creation of an admin-group if it re-uses the same group value but with a different name than an existing group.

Only the admin groups bound to an MPLS interface are advertised area-wide in TE link TLVs and sub-TLVs when the **traffic-engineering** option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

| | |
|-------------------|---|
| Parameters | <p><i>group-name</i> — Specifies the name of the group with up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain</p> <p>value <i>group-value</i> — Specifies the integer value associated with the group. The association of group name and value should be unique within an IP/MPLS domain.</p> <p>Values 0 to 31</p> |
|-------------------|---|

admin-group

| | |
|--------------------|--|
| Syntax | <p>admin-group <i>group-name</i> [<i>group-name...</i>(up to 5 max)]</p> <p>no admin-group <i>group-name</i> [<i>group-name...</i>(up to 5 max)]</p> <p>no admin-group</p> |
| Context | <p>config>router>if>if-attribute</p> <p>config>service>ies>if>if-attribute</p> <p>config>service>vprn>if>if-attribute</p> <p>config>router>mpls>if</p> |
| Description | <p>This command configures the admin group membership of an interface. The user can apply admin groups to an IES, VPRN, network IP, or MPLS interface.</p> |

Each single operation of the **admin-group** command allows a maximum of five (5) groups to be specified at a time. However, a maximum of 32 groups can be added to a given interface through multiple operations. Once an admin group is bound to one or more interface, its value cannot be changed until all bindings are removed.

The configured admin-group membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas.

Only the admin groups bound to an MPLS interface are advertised area-wide in TE link TLVs and sub-TLVs when the **traffic-engineering** option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

The **no** form of this command deletes one or more of the admin-group memberships of an interface. The user can also delete all memberships of an interface by not specifying a group name.

Parameters *group-name* — Specifies the name of the group with up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain.

srlg-group

Syntax **srlg-group** *group-name* **value** *group-value* [**penalty-weight** *penalty-weight*]
no srlg-group *group-name*

Context config>router>if-attribute

Description This command defines a Shared Risk Link Group (SRLG) which can be associated with an IP or MPLS interface.

SRLG is used to tag IP or MPLS interfaces which share a specific fate with the same identifier. For example, an SRLG group identifier could represent all links which use separate fibers but are carried in the same fiber conduit. If the conduit is accidentally cut, all the fiber links are cut which means all interfaces using these fiber links will fail.

The user first configures locally on each router the name and identifier of each SRLG group. A maximum of 1024 SRLGs can be configured per system.

The user then configures the SRLG membership of an interface. The user can apply SRLGs to an IES, VPRN, network IP, or MPLS interface. A maximum of 64 SRLGs can be applied to a given interface.

When SRLGs are applied to MPLS interfaces, CSPF at an LER will exclude the SRLGs of interfaces used by the LSP primary path when computing the path of the secondary path. CSPF at an LER or LSR will also exclude the SRLGs of the outgoing interface of the primary LSP path in the computation of the path of the FRR backup LSP. This provides path disjointness between the primary path and the secondary path or FRR backup path of an LSP.

When SRLGs applied to IES, VPRN, or network IP interfaces, they are evaluated in the route next-hop selection by adding the **srlg-enable** option in a route next-hop policy template applied to an interface or a set of prefixes. For instance, the user can enable the SRLG constraint to select a LFA next-hop for a prefix which avoids all interfaces that share fate with the primary next-hop.

The following provisioning rules are applied to SRLG configuration. The system will reject the creation of a SRLG if it re-uses the same name but with a different group value than an existing group. The system will also reject the creation of an SRLG if it re-uses the same group value but with a different name than an existing group.

Only the SRLGs bound to an MPLS interface are advertised area-wide in TE link TLVs and sub-TLVs when the **traffic-engineering** option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

A user may specify a penalty weight (**penalty-weight**) associated with an SRLG. This controls the likelihood of paths with links sharing SRLG values with a primary path being used by a bypass or detour LSP. The higher the penalty weight, the less desirable it is to use the link with a given SRLG.

| | |
|-------------------|--|
| Parameters | <i>group-name</i> — Specifies the name of the group, up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain. |
| | <i>value group-value</i> — Specifies the integer value associated with the group. The association of group name and value should be unique within an IP/MPLS domain. |
| | Values 0 to 4294967295 |
| | <i>penalty-weight penalty-weight</i> — Specifies the integer value of the penalty weight that is assigned to the SRLG group |
| | Values 0 to 65535 |
| | Default 0 |

srlg-group

| | |
|--------------------|---|
| Syntax | srlg-group <i>group-name</i> [<i>group-name...</i> (up to 5 max)] no srlg-group <i>group-name</i> [<i>group-name...</i> (up to 5 max)] no srlg-group |
| Context | config>router>if>if-attribute config>service>ies>if>if-attribute config>service>vprn>if>if-attribute config>router>mpls>if |
| Description | This command configures the SRLG membership of an interface. The user can apply SRLGs to an IES, VPRN, network IP, or MPLS interface. |

An interface can belong to up to 64 SRLG groups. However, each single operation of the **srlg-group** command allows a maximum of five (5) groups to be specified at a time. Once an SRLG group is bound to one or more interface, its value cannot be changed until all bindings are removed.

The configured SRLG membership will be applied in all levels/areas the interface is participating in. The same interface cannot have different memberships in different levels/areas.

Only the SRLGs bound to an MPLS interface are advertised area-wide in TE link TLVs and sub-TLVs when the **traffic-engineering** option is enabled in IS-IS or OSPF. IES and VPRN interfaces do not have their attributes advertised in TE TLVs.

The **no** form of this command deletes one or more of the SRLG memberships of an interface. The user can also delete all memberships of an interface by not specifying a group name.

Parameters *group-name* — Specifies the name of the group, up to 32 characters. The association of group name and value should be unique within an IP/MPLS domain.

route-next-hop-policy

Syntax **route-next-hop-policy**

Context config>router

Description This command creates the context to configure route next-hop policies.

template

Syntax [no] **template** *template-name*

Context config>router>route-next-hop-policy

Description This command creates a template to configure the attributes of a Loop-Free Alternate (LFA) Shortest Path First (SPF) policy. An LFA SPF policy allows the user to apply specific criteria, such as admin group and SRLG constraints, to the selection of an LFA backup next-hop for a subset of prefixes that resolve to a specific primary next-hop.

The user first creates a route next-hop policy template under the global router context and then applies it to a specific OSPF or IS-IS interface in the global routing instance or in a VPRN instance.

A policy template can be used in both IS-IS and OSPF to apply the specific criteria to prefixes protected by LFA. Each instance of IS-IS or OSPF can apply the same policy template to one or more interface.

The commands within the route next-hop policy template use the **begin-commit-abort** model. The following are the steps to create and modify the template:

To create a template, the user enters the name of the new template directly under the route-next-hop-policy context.

1. To delete a template that is not in use, the user enters the **no** form for the template name under the route-next-hop-policy context.
2. The user enters the editing mode by executing the begin command under the route-next-hop-policy context. The user can then edit and change any number of route next-hop policy templates. However, the parameter value will still be stored temporarily in the template module until the commit is executed under the route-next-hop-policy context. Any temporary parameter changes will be lost if the user enters the abort command before the commit command.
3. The user is allowed to create or delete a template instantly once in the editing mode without the need to enter the commit command. Furthermore, the abort command, if entered, will have no effect on the prior deletion or creation of a template.

Once the commit command is issued, IS-IS or OSPF will re-evaluate the templates and if there are any net changes, it will schedule a new LFA SPF to re-compute the LFA next-hop for the prefixes associated with these templates.

Parameters *template-name* — Specifies the name of the template, up to 32 characters.

include-group

Syntax **include-group** *group-name* [**pref** *pref*]
 no include-group *group-name*

Context config>router>route-next-hop-policy>template

Description This command configures the admin group constraint into the route next-hop policy template.

Each group is entered individually. The **include-group** statement instructs the LFA SPF selection algorithm to pick up a subset of LFA next-hops among the links which belong to one or more of the specified admin groups. A link which does not belong to at least one of the admin-groups is excluded. However, a link can still be selected if it belongs to one of the groups in a include-group statement but also belongs to other groups which are not part of any include-group statement in the route next-hop policy.

The **pref** option is used to provide a relative preference for the admin group to select. A lower preference value means that LFA SPF will first attempt to select a LFA backup next-hop which is a member of the corresponding admin group. If none is found, then the admin group with the next higher preference value is evaluated. If no preference is configured for a given admin group name, then it is supposed to be the least preferred, that is, numerically the highest preference value.

When evaluating multiple **include-group** statements within the same preference, any link which belongs to one or more of the included admin groups can be selected as an LFA next-hop. There is no relative preference based on how many of those included admin groups the link is a member of.

The **exclude-group** statement simply prunes all links belonging to the specified admin group before making the LFA backup next-hop selection for a prefix.

If the same group name is part of both include and exclude statements, the exclude statement will win. In other words, the exclude statement can be viewed as having an implicit preference value of 0.

The admin-group criteria are applied before running the LFA next-hop selection algorithm.

The **no** form deletes the admin group constraint from the route next-hop policy template.

| | |
|-------------------|---|
| Parameters | <i>group-name</i> — Specifies the name of the group, up to 32 characters. |
| | pref <i>pref</i> — An integer specifying the relative preference of a group. |
| | Values 1 to 255 |
| | Default 255 |

exclude-group

| | |
|--------------------|--|
| Syntax | exclude-group <i>group-name</i> no exclude-group <i>group-name</i> |
| Context | config>router>route-next-hop-policy>template |
| Description | <p>This command configures the admin group constraint into the route next-hop policy template.</p> <p>Each group is entered individually. The include-group statement instructs the LFA SPF selection algorithm to pick up a subset of LFA next-hops among the links that belong to one or more of the specified admin groups. A link that does not belong to at least one of the admin-groups is excluded. However, a link can still be selected if it belongs to one of the groups in an include-group statement but also belongs to other groups that are not part of any include-group statement in the route next-hop policy.</p> <p>The pref option is used to provide a relative preference for the admin group to select. A lower preference value means that LFA SPF will first attempt to select an LFA backup next-hop that is a member of the corresponding admin group. If none is found, then the admin group with the next highest preference value is evaluated. If no preference is configured for a given admin group name, then it is supposed to be the least preferred, that is, numerically the highest preference value.</p> <p>When evaluating multiple include-group statements within the same preference, any link that belongs to one or more of the included admin groups can be selected as an LFA next-hop. There is no relative preference based on how many of those included admin groups the link is a member of.</p> <p>The exclude-group statement simply prunes all links belonging to the specified admin group before making the LFA backup next-hop selection for a prefix.</p> |

If the same group name is part of both include and exclude statements, the exclude statement will win. In other words, the exclude statement can be viewed as having an implicit preference value of zero (0).

The admin-group criteria are applied before running the LFA next-hop selection algorithm.

The **no** form deletes the admin group constraint from the route next-hop policy template.

Parameters *group-name* — Specifies the name of the group, up to 32 characters.

srlg-enable

Syntax **[no] srlg-enable**

Context config>router>route-next-hop-policy>template

Description This command configures the SRLG constraint into the route next-hop policy template.

When this command is applied to a prefix, the LFA SPF will attempt to select an LFA next-hop, among the computed ones, which uses an outgoing interface that does not participate in any of the SLRGs of the outgoing interface used by the primary next-hop.

The SRLG criterion is applied before running the LFA next-hop selection algorithm.

The **no** form deletes the SRLG constraint from the route next-hop policy template.

Default no srlg-enable

protection-type

Syntax **protection-type {link | node}**
no protection-type

Context config>router>route-next-hop-policy>template

Description This command configures the protection type constraint into the route next-hop policy template.

The user can select if link protection or node protection is preferred in the selection of an LFA next-hop for all IP prefixes and LDP FEC prefixes to which a route next-hop policy template is applied. The default in SR OS implementation is node protection. The implementation will fall back to the other type if no LFA next-hop of the preferred type is found.

When the route next-hop policy template is applied to an IP interface, all prefixes using this interface as a primary next-hop will follow the protection type preference specified in the template.

The **no** form deletes the protection type constraint from the route next-hop policy template.

| | |
|------------|---|
| Default | protection-type node |
| Parameters | <i>{link node}</i> — Specifies the two possible values for the protection type. |
| Default | node |

nh-type

| | |
|-------------|---|
| Syntax | nh-type {ip tunnel} no nh-type |
| Context | config>router>route-next-hop-policy>template |
| Description | <p>This command configures the next-hop type constraint into the route next-hop policy template.</p> <p>The user can select if tunnel backup next-hop or IP backup next-hop is preferred. The default in SR OS implementation is to prefer IP next-hop over tunnel next-hop. The implementation will fall back to the other type if no LFA next-hop of the preferred type is found.</p> <p>When the route next-hop policy template is applied to an IP interface, all prefixes using this interface as a primary next-hop will follow the next-hop type preference specified in the template.</p> <p>The no form deletes the next-hop type constraint from the route next-hop policy template.</p> |
| Default | nh-type ip |
| Parameters | <i>{ip tunnel}</i> — Specifies the two possible values for the next-hop type. |
| Default | ip |

2.12.2.4.1 Router Interface Filter Commands

egress

| | |
|-------------|---|
| Syntax | egress |
| Context | config>router>if |
| Description | This command enables access to the context to configure egress network filter policies for the IP interface. If an egress filter is not defined, no filtering is performed. |

ingress

| | |
|--------------------|---|
| Syntax | ingress |
| Context | config>router>if |
| Description | This command enables access to the context to configure ingress network filter policies for the IP interface. If an ingress filter is not defined, no filtering is performed. |

filter

| | |
|--------------------|--|
| Syntax | filter ip <i>ip-filter-id</i> filter ipv6 <i>ipv6-filter-id</i> no filter [ip <i>ip-filter-ip</i>] [ipv6 <i>ipv6-filter-id</i>] |
| Context | config>router>if>ingress config>router>if>egress |
| Description | <p>This command associates an IP filter policy with an IP interface.</p> <p>Filter policies control packet forwarding and dropping based on IP match criteria.</p> <p>The <i>ip-filter-id</i> must have been preconfigured before this filter command is executed. If the filter ID does not exist, an error occurs.</p> <p>Only one filter ID can be specified.</p> <p>The no form of the command removes the filter policy association with the IP interface.</p> |
| Default | No filter is specified. |
| Parameters | <p><i>ip ip-filter-id</i> — The filter name acts as the ID for the IP filter policy expressed as a decimal integer. The filter policy must already exist within the config>filter>ip context.</p> <p>Values 1 to 16384</p> <p><i>ipv6 ipv6-filter-id</i> — The filter name acts as the ID for the IPv6 filter policy expressed as a decimal integer. The filter policy must already exist within the config>filter>ipv6 context. This parameter only applies to the 7750 SR and 7950 XRS.</p> <p>Values 1 to 65535</p> |

2.12.2.4.2 Router Interface ICMP Commands

hold-time

| | |
|--------------------|---|
| Syntax | hold-time |
| Context | <pre> config>router>if config>service>ies>if config>service>ies>subscriber-interface config>service>ies>redundant-interface config>service>vprn>if config>service>vprn>network-interface config>service>vprn>subscriber-interface config>service>vprn>redundant-interface config>service>vpls>if </pre> |
| Description | <p>This command creates the CLI context to configure interface level hold-up and hold-down timers for the associated IP interface.</p> <p>The up timer controls a delay for the associated IPv4 or IPv6 interface so that the system will delay the deactivation of the associated interface for the specified amount of time.</p> <p>The down timer controls a delay for the associated IPv4 or IPv6 interface so that the system will delay the activation of the associated interface for the specified amount of time</p> |

up

| | |
|--------------------|--|
| Syntax | <pre> up ip seconds no up ip up ipv6 seconds no up ipv6 </pre> |
| Context | <pre> config>router>if>hold-time config>service>ies>if>hold-time config>service>ies>sub-if>hold-time config>service>ies>red-if>hold-time config>service>vprn>if>hold-time config>service>vprn>nw-if>hold-time config>service>vprn>sub-if>hold-time config>service>vprn>red-if>hold-time config>service>vpls>if>hold-time </pre> |
| Description | <p>This command will cause a delay in the deactivation of the associated IP interface by the specified number of seconds. The delay is invoked whenever the system attempts to bring the associated IP interface down.</p> |

The **no** form of the command removes the command from the active configuration and removes the delay in deactivating the associated IP interface. If the configuration is removed during a delay period, the currently running delay will continue until it expires.

Default no up ip

Parameters *seconds* — The time delay, in seconds, to make the interface operational.

Values 1 to 1200

down

Syntax **down ip** *seconds* [**init-only**]
no up ip
up ipv6 *seconds* [**init-only**]
no up ipv6

Context config>router>if>hold-time
config>service>ies>if>hold-time
config>service>ies>sub-if>hold-time
config>service>ies>red-if>hold-time
config>service>vprn>if>hold-time
config>service>vprn>nw-if>hold-time
config>service>vprn>sub-if>hold-time
config>service>vprn>red-if>hold-time
config>service>vpls>if>hold-time

Description This command will cause a delay in the activation of the associated IP interface by the specified number of seconds. The delay is invoked whenever the system attempts to bring the associated IP interface up, unless the **init-only** option is configured. If the **init-only** option is configured, the delay is only applied when the IP interface is first configured or after a system reboot.

The **no** form of the command removes the command from the active configuration and removes the delay in activating the associated IP interface. If the configuration is removed during a delay period, the currently running delay will continue until it completes.

Default no down ip

Parameters *seconds* — The time delay, in seconds, to make the interface operational.

Values 1 to 1200

init-only — Specifies that the **down** delay is only applied when the interface is configured or after a reboot.

Values 1 to 1200

icmp

| | |
|--------------------|--|
| Syntax | icmp |
| Context | config>router>if |
| Description | This command enables access to the context to configure Internet Control Message Protocol (ICMP) parameters on a network IP interface. ICMP is a message control and error reporting protocol that also provides information relevant to IP packet processing. |

mask-reply

| | |
|--------------------|---|
| Syntax | [no] mask-reply |
| Context | config>router>if>icmp |
| Description | <p>This command enables responses to ICMP mask requests on the router interface.</p> <p>If a local node sends an ICMP mask request to the router interface, the mask-reply command configures the router interface to reply to the request.</p> <p>The no form of the command disables replies to ICMP mask requests on the router interface.</p> |
| Default | mask-reply — Replies to ICMP mask requests. |

param-problem

| | |
|--------------------|---|
| Syntax | param-problem [<i>number seconds</i>] no param-problem |
| Context | config>router>if>icmp config>router>if>icmp6 |
| Description | <p>This command specifies whether parameter-problem ICMP messages should be sent. When enabled, parameter-problem ICMP messages are generated by this interface.</p> <p>The no form of the command disables the sending of parameter-problem ICMP messages.</p> |
| Parameters | <p><i>number</i> — Specifies the number of parameter-problem ICMP messages to send in the time frame specified by the <i>seconds</i> parameter.</p> <p>Values 10 to 1000</p> <p>Default 100</p> <p><i>seconds</i> — Specifies the time frame, in seconds, that is used to limit the number of parameter-problem ICMP messages issued.</p> <p>Values 1 to 60</p> <p>Default 10</p> |

redirects

| | |
|--------------------|--|
| Syntax | redirects [<i>number seconds</i>] no redirects |
| Context | config>router>if>icmp |
| Description | <p>This command enables and configures the rate for ICMP redirect messages issued on the router interface.</p> <p>When routes are not optimal on this router, and another router on the same subnetwork has a better route, the router can issue an ICMP redirect to alert the sending node that a better route is available.</p> <p>The redirects command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects are issued can be controlled with the optional <i>number</i> and <i>time</i> parameters by indicating the maximum number of redirect messages that can be issued on the interface for a given time interval.</p> <p>By default, generation of ICMP redirect messages is enabled at a maximum rate of 100 per 10 second time interval.</p> <p>The no form of the command disables the generation of ICMP redirects on the router interface.</p> |
| Default | redirects 100 10 — Maximum of 100 redirect messages in 10 seconds. |
| Parameters | <p><i>number</i> — The maximum number of ICMP redirect messages to send, expressed as a decimal integer. This parameter must be specified with the <i>time</i> parameter.</p> <p>Values 10 to 1000</p> <p><i>seconds</i> — The time frame, in seconds, used to limit the <i>number</i> of ICMP redirect messages that can be issued, expressed as a decimal integer.</p> <p>Values 1 to 60</p> |

ttl-expired

| | |
|--------------------|---|
| Syntax | ttl-expired [<i>number seconds</i>] no ttl-expired |
| Context | config>router>if>icmp |
| Description | <p>This command configures the rate that Internet Control Message Protocol (ICMP) Time To Live (TTL) expired messages are issued by the IP interface.</p> <p>By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.</p> <p>The no form of the command disables the generation of TTL expired messages.</p> |

| | |
|-------------------|--|
| Default | ttl-expired 100 10 — Maximum of 100 TTL expired message in 10 seconds. |
| Parameters | <p><i>number</i> — The maximum number of ICMP TTL expired messages to send, expressed as a decimal integer. The <i>seconds</i> parameter must also be specified.</p> <p>Values 10 to 2000</p> <p><i>seconds</i> — The time frame, in seconds, used to limit the <i>number</i> of ICMP TTL expired messages that can be issued, expressed as a decimal integer.</p> <p>Values 1 to 60</p> |

unreachables

| | |
|--------------------|--|
| Syntax | unreachables [<i>number seconds</i>] no unreachable s |
| Context | config>router>if>icmp |
| Description | <p>This command enables and configures the rate for ICMP host and network destination unreachable messages issued on the router interface.</p> <p>The unreachables command enables the generation of ICMP destination unreachables on the router interface. The rate at which ICMP unreachables is issued can be controlled with the optional <i>number</i> and <i>seconds</i> parameters by indicating the maximum number of destination unreachable messages that can be issued on the interface for a given time interval.</p> <p>By default, generation of ICMP destination unreachables messages is enabled at a maximum rate of 100 per 10 second time interval.</p> <p>The no form of the command disables the generation of ICMP destination unreachables on the router interface.</p> |
| Default | unreachables 100 10 — Maximum of 100 unreachable messages in 10 seconds. |
| Parameters | <p><i>number</i> — The maximum number of ICMP unreachable messages to send, expressed as a decimal integer. The <i>seconds</i> parameter must also be specified.</p> <p>Values 10 to 2000</p> <p><i>seconds</i> — The time frame, in seconds, used to limit the <i>number</i> of ICMP unreachable messages that can be issued, expressed as a decimal integer.</p> |

2.12.2.4.3 Router Interface IPv6 Commands

ipv6

| | |
|--------------------|---|
| Syntax | [no] ipv6 |
| Context | config>router>if |
| Description | This command configures IPv6 for a router interface. The no form of the command disables IPv6 on the interface. |
| Default | not enabled |

address

| | |
|--------------------|--|
| Syntax | address { <i>ipv6-address/prefix-length</i> } [eui-64] [preferred] [track-srrp <i>srrp-instance</i>] [modifier <i>cga-modifier</i>] no address { <i>ipv6-address/prefix-length</i> } |
| Context | config>router>if>ipv6 |
| Description | This command assigns an IPv6 address to the interface. Up to 16 total primary and secondary IPv4 and IPv6 addresses can be assigned to network interfaces, and up to 256 to access interfaces. |



Caution: Configurations must not exceed 16 IPv6 addresses when IPSec, GRE, L2TPv3, or IP in IP protocols are active on an access interface.

| | |
|-------------------|--|
| Parameters | <i>ipv6-address/prefix-length</i> — Specifies the IPv6 address on the interface. |
|-------------------|--|

Values

| | | |
|----------------------|---------------|---|
| ipv6-address/prefix: | ipv6-address | x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x [0 to FFFF]H d [0 to 255]D |
| | prefix-length | 1 to 128 |

eui-64 — When the **eui-64** keyword is specified, a complete IPv6 address from the supplied prefix and 64-bit interface identifier is formed. The 64-bit interface identifier is derived from MAC address on Ethernet interfaces. For interfaces without a MAC address, for example POS interfaces, the Base MAC address of the chassis should be used.

preferred — Specifies that the IPv6 address is the preferred IPv6 address for this interface. The preferred address is an address assigned to an interface whose use by upper layer protocols is unrestricted. Preferred addresses may be used as the source (or destination) address of packets sent from (or to) the interface. Preferred addresses do not go through the DAD process.

track-srrp — Specifies the SRRP instance ID that this interface route needs to track.

srrp-instance — Indicates the unique identifier of the tracked SRRP instance.

Values 1 to 4294967295

modifier *cga-modifier* — Sets the modifier for cryptographically-assigned addresses.

Values 0x0..0xFFFFFFFF...(32 hex nibbles)

dad-disable

| | |
|--------------------|---|
| Syntax | [no] dad-disable |
| Context | config>router>if>ipv6 |
| Description | This command disables duplicate address detection (DAD) on a per-interface basis. This prevents the router from performing a DAD check on the interface. All IPv6 addresses of an interface with DAD disabled, immediately enter a preferred state, without checking for uniqueness on the interface. This is useful for interfaces which enter a looped state during troubleshooting and operationally disable themselves when the loop is detected, requiring manual intervention to clear the DAD violation. The no form of the command turns off dad-disable on the interface. |
| Default | not enabled |

icmp6

| | |
|--------------------|--|
| Syntax | icmp6 |
| Context | config>router>if>ipv6 |
| Description | This command enables the context to configure ICMPv6 parameters for the interface. |

packet-too-big

| | |
|--------------------|---|
| Syntax | packet-too-big [<i>number seconds</i>] no packet-too-big |
| Context | config>router>if>ipv6>icmp6 |
| Description | This command configures the rate for ICMPv6 packet-too-big messages. |

| | |
|-------------------|--|
| Parameters | <i>number</i> — Limits the number of packet-too-big messages issued per time frame specified in the <i>seconds</i> parameter. Values 10 to 1000 <i>seconds</i> — Determines the time frame, in seconds, that is used to limit the number of packet-too-big messages issued per time frame. Values 1 to 60 |
|-------------------|--|

redirects

| | |
|--------------------|--|
| Syntax | redirects [<i>number seconds</i>] no redirects |
| Context | config>router>if>ipv6>icmp6 |
| Description | This command configures the rate for ICMPv6 redirect messages. When configured, ICMPv6 redirects are generated when routes are not optimal on the router and another router on the same subnetwork has a better route to alert that node that a better route is available. The no form of the command disables ICMPv6 redirects. |
| Default | 100 10 (when IPv6 is enabled on the interface) |
| Parameters | <i>number</i> — Limits the number of redirects issued per the time frame specified in <i>seconds</i> parameter. Values 10 to 1000 <i>seconds</i> — Determines the time frame, in seconds, that is used to limit the number of redirects issued per time frame. Values 1 to 60 |

time-exceeded

| | |
|--------------------|--|
| Syntax | time-exceeded [<i>number seconds</i>] no time-exceeded |
| Context | config>router>if>ipv6>icmp6 |
| Description | This command configures rate for ICMPv6 time-exceeded messages. |
| Parameters | <i>number</i> — Limits the number of time-exceeded messages issued per the time frame specified in <i>seconds</i> parameter. Values 10 to 2000 |

seconds — Determines the time frame, in seconds, that is used to limit the number of time-exceeded messages issued per time frame.

Values 1 to 60

unreachables

| | |
|--------------------|---|
| Syntax | unreachables [<i>number seconds</i>] no unreachables |
| Context | config>router>if>ipv6>icmp6 |
| Description | This command configures the rate for ICMPv6 unreachable messages. When enabled, ICMPv6 host and network unreachable messages are generated by this interface. The no form of the command disables the generation of ICMPv6 host and network unreachable messages by this interface. |
| Default | 100 10 (when IPv6 is enabled on the interface) |
| Parameters | <i>number</i> — Determines the number destination unreachable ICMPv6 messages to issue in the time frame specified in <i>seconds</i> parameter. Values 10 to 2000 <i>seconds</i> — Sets the time frame, in seconds, to limit the number of destination unreachable ICMPv6 messages issued per time frame. Values 1 to 60 |

link-local-address

| | |
|--------------------|--|
| Syntax | link-local-address <i>ipv6-address</i> [dad-disable] |
| Context | config>router>if>ipv6 config>service>ies>if>ipv6 config>service>vprn>if>ipv6 |
| Description | This command configures the IPv6 link local address. The no form of the command removes the configured link local address, and the router automatically generates a default link local address. Removing a manually configured link local address may impact routing protocols or static routes that have a dependency on that address. It is not recommended to remove a link local address when there are active IPv6 subscriber hosts on an IES or VPRN interface. |
| Parameters | dad-disable — Disables Duplicate Address Detection (DAD) and sets the address to preferred, even if there is a duplicated address. |

local-proxy-nd

| | |
|--------------------|---|
| Syntax | [no] local-proxy-nd |
| Context | config>router>if>ipv6 |
| Description | This command enables local proxy neighbor discovery on the interface. The no form of the command disables local proxy neighbor discovery. |

neighbor

| | |
|-----------------------|---|
| Syntax | neighbor <i>[ipv6-address]</i> <i>[mac-address]</i> no neighbor <i>[ipv6-address]</i> |
| Context | config>router>if>ipv6 |
| Description | This command configures an IPv6-to-MAC address mapping on the interface. Use this command if a directly attached IPv6 node does not support ICMPv6 neighbor discovery, or for some reason, a static address must be used. This command can only be used on Ethernet media. The <i>ipv6-address</i> must be on the subnet that was configured from the IPv6 address command or a link-local address. |
| Parameters | <i>ipv6-address</i> — The IPv6 address assigned to a router interface. |
| | Values |
| <i>ipv6-address</i> : | x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D |
| | <i>mac-address</i> — Specifies the MAC address for the neighbor in the form of xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx. |

neighbor-limit

| | |
|--------------------|---|
| Syntax | neighbor-limit <i>limit</i> [log-only] [threshold percent] no neighbor-limit |
| Context | config>router>if>ipv6 |
| Description | This command configures the maximum amount of dynamic IPv6 neighbor entries that can be learned on an IP interface. |

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

The **no** form of the command removes the neighbor-limit.

| | |
|-------------------|--|
| Default | 90 percent |
| Parameters | <p>log-only — Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, entries above the limit will be learned.</p> <p><i>percent</i> — The threshold value (as a percentage) that triggers a warning message to be sent.</p> <p>Values 0 to 100</p> <p><i>limit</i> — The number of entries that can be learned on an IP interface expressed as a decimal integer. If the limit is set to 0, dynamic neighbor learning is disabled and no dynamic neighbor entries are learned.</p> <p>Values 0 to 102400</p> |

proxy-nd-policy

| | |
|--------------------|---|
| Syntax | proxy-nd-policy <i>policy-name</i> [<i>policy-name...</i> (up to 5 max)] no proxy-nd-policy |
| Context | config>router>if>ipv6 |
| Description | This command configure a proxy neighbor discovery policy for the interface. |
| Parameters | <i>policy-name</i> — The neighbor discovery policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. The specified policy names must already be defined. |

2.12.2.4.4 Router Interface DHCP Commands

dhcp

| | |
|--------------------|--|
| Syntax | dhcp |
| Context | config>router>if |
| Description | This command enables the context to configure DHCP parameters. |

gi-address

| | |
|--------------------|---|
| Syntax | gi-address ip-address [src-ip-addr] no gi-address |
| Context | config>router>if>dhcp |
| Description | This command configures the gateway interface address for the DHCP relay. The GI address is needed, when the router functions as a DHCP relay, to distinguish between the different subscriber interfaces and potentially between the group interfaces defined. |
| Default | no gi-address |
| Parameters | <i>ip-address</i> — Specifies the host IP address to be used for DHCP relay packets. <i>src-ip-address</i> — Specifies the source IP address to be used for DHCP relay packets. |

option

| | |
|--------------------|---|
| Syntax | [no] option |
| Context | config>router>if>dhcp |
| Description | This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options. The no form of this command returns the system to the default. |
| Default | no option |

action

| | |
|--------------------|---|
| Syntax | action {replace drop keep} no action |
| Context | config>router>if>dhcp>option |
| Description | This command configures the processing required when the SR-Series router receives a DHCP request that already has a Relay Agent Information Option (Option 82) field in the packet. The no form of this command returns the system to the default value. |
| Default | Per RFC 3046, DHCP Relay Agent Information Option, section 2.1.1, Reforwarded DHCP requests, the default is to keep the existing information intact. The exception to this is if the GI address of the received packet is the same as the ingress address on the router. In that case the packet is dropped and an error is logged. |

- Parameters**
- replace** — In the upstream direction (from the user), the existing Option 82 field is replaced with the Option 82 field from the router. In the downstream direction (toward the user) the Option 82 field is stripped (in accordance with RFC 3046).
 - drop** — The packet is dropped, and an error is logged.
 - keep** — The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is sent on toward the client.
The behavior is slightly different in case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router will insert his own VSO into the Option 82 field. This will only be done when the incoming message has already an Option 82 field.
If no Option 82 field is present, the router will not create the Option 82 field. In this in that case, no VSO will be added to the message.

circuit-id

- Syntax** **circuit-id [ascii-tuple | ifindex | sap-id | vlan-ascii-tuple]**
no circuit-id
- Context** config>router>if>dhcp>option
- Description** When enabled, the router sends the interface index (If Index) in the **circuit-id** suboption of the DHCP packet. The If Index of a router interface can be displayed using the command **show>router>if>detail**. This option specifies data that must be unique to the router that is relaying the circuit.

If disabled, the **circuit-id** suboption of the DHCP packet will be left empty.

The **no** form of this command returns the system to the default.
- Default** circuit-id ascii-tuple
- Parameters**
- ascii-tuple** — Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by “|”.
 - ifindex** — Specifies that the interface index will be used. The If Index of a router interface can be displayed using the command **show>router>if>detail**.
 - sap-id** — Specifies that the SAP ID will be used.
 - vlan-ascii-tuple** — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and qinq ports only. Therefore, when the Option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

remote-id

| | |
|--------------------|---|
| Syntax | remote-id [mac string <i>string</i>] no remote-id |
| Context | config>router>if>dhcp>option |
| Description | <p>When enabled, the router sends the MAC address of the remote end (typically the DHCP client) in the remote-id suboption of the DHCP packet. This command identifies the host at the other end of the circuit. If disabled, the remote-id suboption of the DHCP packet will be left empty.</p> <p>The no form of this command returns the system to the default.</p> |
| Default | no remote-id |
| Parameters | <p>mac — This keyword specifies the MAC address of the remote end is encoded in the suboption.</p> <p>string <i>string</i> — Specifies the remote-id.</p> |

vendor-specific-option

| | |
|--------------------|---|
| Syntax | [no] vendor-specific-option |
| Context | config>router>if>dhcp>option |
| Description | This command configures the Nokia vendor specific suboption of the DHCP relay packet. |

client-mac-address

| | |
|--------------------|---|
| Syntax | [no] client-mac-address |
| Context | config>router>if>dhcp>option |
| Description | <p>This command enables the sending of the MAC address in the Nokia vendor specific suboption of the DHCP relay packet.</p> <p>The no form of the command disables the sending of the MAC address in the Nokia vendor specific suboption of the DHCP relay packet.</p> |
| Default | no client-mac-address |

pool-name

| | |
|---------------|--------------------------------|
| Syntax | [no] pool-name |
|---------------|--------------------------------|

Context config>router>if>dhcp>option>vendor-specific-option

Description This command enables the sending of the pool name in the Nokia vendor-specific suboption of the DHCP relay packet.

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

Default no pool-name

port-id

Syntax [no] port-id

Context config>router>if>dhcp>option>vendor-specific-option

Description This command enables sending of the port-id in the Nokia vendor specific suboption of the DHCP relay packet

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

Default no port-id

service-id

Syntax [no] service-id

Context config>router>if>dhcp>option>vendor-specific-option

Description This command enables the sending of the service ID in the Nokia vendor specific suboption of the DHCP relay packet.

The **no** form of the command disables the sending of the service ID in the Nokia vendor specific suboption of the DHCP relay packet.

Default no service-id

string

Syntax [no] string *text*

Context config>router>if>dhcp>option>vendor-specific-option

Description This command specifies the vendor specific suboption string of the DHCP relay packet.

The **no** form of the command returns the default value.

Default no string

| | |
|-------------------|--|
| Parameters | <i>text</i> — The string can be any combination of ASCII characters up to 32 characters in length. If spaces are used in the string, enclose the entire string in quotation marks (" "). |
|-------------------|--|

system-id

| | |
|--------------------|---|
| Syntax | [no] system-id |
| Context | config>router>if>dhcp>option>vendor-specific-option |
| Description | This command specifies whether the system-id is encoded in the Nokia vendor specific sub-option of Option 82. |
| Default | no system-id |

relay-plain-bootp

| | |
|--------------------|---|
| Syntax | [no] relay-plain-bootp |
| Context | config>router>if>dhcp |
| Description | <p>This command enables the relaying of plain BOOTP packets.</p> <p>The no form of the command disables the relaying of plain BOOTP packets.</p> |
| Default | no relay-plain-bootp |

server

| | |
|--------------------|---|
| Syntax | server server1 [server2...(up to 8 max)] |
| Context | config>router>if>dhcp |
| Description | <p>This command specifies a list of servers where requests will be forwarded. The list of servers can entered as either IP addresses or fully qualified domain names. There must be at least one server specified for DHCP relay to work. If there are multiple servers then the request is forwarded to all of the servers in the list. There can be a maximum of 8 DHCP servers configured.</p> <p>The flood command is applicable only in the VPLS case. There is a scenario with VPLS where the VPLS node only wants to add Option 82 information to the DHCP request to provider per-subscriber information, but it does not do full DHCP relay. In this case, the server is set to "flood". This means the DHCP request is still a broadcast and is sent through the VPLS domain. A node running at L3 further upstream then can perform the full L3 DHCP relay function.</p> |
| Default | no server |

Parameters *server* — Specifies the DHCP server IP address.

trusted

Syntax **[no] trusted**

Context config>router>if>dhcp

Description According to RFC 3046, **DHCP Relay Agent Information Option**, a DHCP request where the GI address is 0.0.0.0 and which contains an Option 82 field in the packet, should be discarded, unless it arrives on a "trusted" circuit.

If trusted mode is enabled on an IP interface, the relay agent (the SR-Series) will modify the request's GI address to be equal to the ingress interface and forward the request.

This behavior only applies when the action in the Relay Agent Information Option is "keep". In the case where the Option 82 field is being replaced by the relay agent (action = "replace"), the original Option 82 information is lost anyway, and there is no reason for enabling the trusted option.

The **no** form of this command returns the system to the default.

Default no trusted

python-policy

Syntax **python-policy** *name*
 no python-policy

Context config>router>if>dhcp

Description This command specifies a python policy. Python policies are configured in the **config>python> python-policy** *name* context.

Default no python-policy

Parameters *name* — Specifies the name of an existing python script up to 32 characters in length.

2.12.2.5 Router Interface Encryption Commands

The following commands are only applicable to platforms that support network group encryption (NGE).

group-encryption

| | |
|--------------------|---|
| Syntax | [no] group-encryption |
| Context | config>router>interface |
| Description | <p>This command enables NGE on the router interface. When NGE is enabled on the interface, all received Layer 3 packets that have the protocol ID configured as ESP are considered to be NGE packets and must be encrypted using a valid set of keys from any preconfigured key group on the system.</p> <p>The no form of the command disables NGE on the interface. NGE cannot be disabled unless all key groups and IP exception filters are removed.</p> |
| Default | no group-encryption |

encryption-keygroup

| | |
|--------------------|---|
| Syntax | encryption-keygroup <i>keygroup-id</i> direction {inbound outbound} no encryption-keygroup direction {inbound outbound} |
| Context | config>router>if>group-encryption |
| Description | <p>This command is used to bind a key group to a router interface for inbound or outbound packet processing. When configured in the outbound direction, packets egressing the router use the active-outbound-sa associated with the configured key group. When configured in the inbound direction, received packets must be encrypted using one of the valid security associations configured for the key group.</p> <p>The no form of the command removes the key group from the router interface in the specified direction.</p> |
| Default | no encryption-keygroup direction inbound no encryption-keygroup direction outbound |
| Parameters | <p><i>keygroup-id</i> — The ID number of the key group being configured.</p> <p>Values 1 to 127, <i>keygroup-name</i> (64 characters maximum)</p> <p>inbound — Binds the key group in the inbound direction.</p> <p>outbound — Binds the key group in the outbound direction.</p> |

ip-exception

| | |
|----------------|--|
| Syntax | ip-exception <i>filter-id</i> direction {inbound outbound} no ip-exception direction {inbound outbound} |
| Context | config>router>if>group-encryption |

| | |
|--------------------|--|
| Description | <p>This command associates an IP exception filter policy with an NGE-enabled router interface to allow packets matching the exception criteria to transit the NGE domain as clear text.</p> <p>When an exception filter is added for inbound traffic, packets matching the criteria in the IP exception filter policy are allowed to be received in clear text even if an inbound key group is configured. If no inbound key group is configured, then associated inbound IP exception filter policies will be ignored.</p> <p>When an exception filter is added for outbound traffic, packets matching the criteria in the IP exception filter policy are not encrypted when sent out of the router interface even if an outbound key group is configured. If no outbound key group is configured, then associated outbound IP exception filter policies will be ignored.</p> <p>The no form of the command removes the IP exception filter policy from the specified direction.</p> |
| Default | <p>no ip-exception direction inbound</p> <p>no ip-exception direction outbound</p> |
| Parameters | <p><i>filter-id</i> — Specifies the IP exception filter policy. The IP exception ID or exception name must have already been created.</p> <p>Values 1 to 6553, <i>filter-name</i> (64 characters maximum)</p> <p>inbound — Binds the exception filter policy in the inbound direction.</p> <p>outbound — Binds the exception filter policy in the outbound direction.</p> |

2.12.2.6 Router Advertisement Commands

router-advertisement

| | |
|--------------------|--|
| Syntax | [no] router-advertisement |
| Context | config>router |
| Description | <p>This command configures router advertisement properties. By default, it is disabled for all IPv6 enabled interfaces.</p> <p>The no form of the command disables all IPv6 interface. However, the no interface interface-name command disables a specific interface.</p> |
| Default | disabled |

dns-options

| | |
|---------------|-------------------------|
| Syntax | [no] dns-options |
|---------------|-------------------------|

| | |
|--------------------|---|
| Context | config>router>router-advert config>router>router-advert>if |
| Description | <p>This command enables the context for configuration of DNS information for Stateless Address Auto-Configuration (SLAAC) hosts. When specified at the router-advertisement level in the routing context, this command allows configuration of service-wide parameters. These can then be inherited at the interface level by specifying the config>router>router-advert>if>dns-options>include-dns command.</p> <p>The no form of the command disables configuration of DNS information for Stateless Address Auto-Configuration (SLAAC) hosts.</p> |
| Default | disabled |

servers

| | |
|--------------------|--|
| Syntax | server <i>ipv6-address</i> no server |
| Context | config>router>router-advert>dns-options config>router>router-advert>if>dns-options |
| Description | This command specifies the IPv6 DNS servers to include in the RDNSS option in Router Advertisements. When specified at the router advertisement level this applies to all interfaces that have include-dns enabled, unless the interfaces have more specific dns-options configured. |
| Parameters | <i>ipv6-address</i> — Specifies the IPv6 address of the DNS servers, up to 4 max. Specified as eight 16-bit hexadecimal pieces. |

include-dns

| | |
|--------------------|--|
| Syntax | [no] include-dns |
| Context | config>router>router-advert>if>dns-options |
| Description | <p>This command enables the Recursive DNS Server (RDNSS) Option in router advertisements. This must be enabled for each interface on which the RDNSS option is required in router advertisement messages.</p> <p>The no form of the command disables the RDNSS option in router advertisements.</p> |
| Default | disabled |

rdnss-lifetime

| | |
|---------------|--|
| Syntax | rdnss-lifetime { <i>seconds</i> infinite } |
|---------------|--|

no rdns-lifetime

| | |
|--------------------|---|
| Context | config>router>router-advert>dns-options config>router>router-advert>if>dns-options |
| Description | This command specifies the maximum time that the RDNS address may be used for name resolution by the client. The RDNS Lifetime must be no more than twice MaxRtrAdvLifetime with a maximum of 3600 seconds. |
| Default | rdns-lifetime infinite |
| Parameters | <i>infinite</i> — Specifies an infinite RDNS lifetime. <i>seconds</i> — Specifies the time in seconds. Values 4 to 3600 |

interface

| | |
|--------------------|---|
| Syntax | [no] interface <i>ip-int-name</i> |
| Context | config>router>router-advert |
| Description | This command configures router advertisement properties on a specific interface. The interface must already exist in the config>router>if context. |
| Default | No interfaces are configured by default. |
| Parameters | <i>ip-int-name</i> — Specifies the interface name. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. |

current-hop-limit

| | |
|--------------------|---|
| Syntax | current-hop-limit <i>number</i> no current-hop-limit |
| Context | config>router>router-advert>if |
| Description | This command configures the current-hop-limit in the router advertisement messages. It informs the nodes on the subnet about the hop-limit when originating IPv6 packets. |
| Default | current-hop-limit 64 |
| Parameters | <i>number</i> — Specifies the hop limit. Values 0 to 255. A value of zero means there is an unspecified number of hops. |

managed-configuration

| | |
|--------------------|--|
| Syntax | [no] managed-configuration |
| Context | config>router>router-advert>if |
| Description | This command sets the managed address configuration flag. This flag indicates that DHCPv6 is available for address configuration in addition to any address autoconfigured using stateless address autoconfiguration. See RFC 3315, <i>Dynamic Host Configuration Protocol (DHCP) for IPv6</i> . |
| Default | no managed-configuration |

max-advertisement-interval

| | |
|--------------------|--|
| Syntax | [no] max-advertisement-interval <i>seconds</i> |
| Context | config>router>router-advert>if |
| Description | This command configures the maximum interval between sending router advertisement messages. |
| Default | max-advertisement-interval 600 |
| Parameters | <i>seconds</i> — Specifies the maximum interval in seconds between sending router advertisement messages. Values 4 to 1800 |

min-advertisement-interval

| | |
|--------------------|--|
| Syntax | [no] min-advertisement-interval <i>seconds</i> |
| Context | config>router>router-advert>if |
| Description | This command configures the minimum interval between sending ICMPv6 neighbor discovery router advertisement messages. |
| Default | min-advertisement-interval 200 |
| Parameters | <i>seconds</i> — Specifies the minimum interval in seconds between sending ICMPv6 neighbor discovery router advertisement messages. Values 3 to 1350 |

mtu

| | |
|---------------|----------------------------------|
| Syntax | [no] mtu <i>mtu-bytes</i> |
|---------------|----------------------------------|

| | |
|--------------------|--|
| Context | config>router>router-advert>if |
| Description | This command configures the MTU for the nodes to use to send packets on the link. |
| Default | no mtu — The MTU option is not sent in the router advertisement messages. |
| Parameters | <i>mtu-bytes</i> — Specifies the MTU for the nodes to use to send packets on the link. |
| Values | 1280 to 9212 |

other-stateful-configuration

| | |
|--------------------|---|
| Syntax | [no] other-stateful-configuration |
| Context | config>router>router-advert>if |
| Description | This command sets the "Other configuration" flag. This flag indicates that DHCPv6lite is available for autoconfiguration of other (non-address) information such as DNS-related information or information about other servers in the network. See RFC 3736, <i>Stateless Dynamic Host Configuration Protocol (DHCP) for IPv6</i> . |
| Default | no other-stateful-configuration |

prefix

| | |
|--------------------|--|
| Syntax | [no] prefix [<i>ipv6-prefix</i> / <i>prefix-length</i>] |
| Context | config>router>router-advert>if |
| Description | This command configures an IPv6 prefix in the router advertisement messages. To support multiple IPv6 prefixes, use multiple prefix statements. No prefix is advertised until explicitly configured using prefix statements. |
| Parameters | <i>ip-prefix</i> — The IP prefix for prefix list entry in dotted decimal notation. |
| Values | |

| | |
|--------------------|-------------------------------------|
| ipv6-prefix | x:x:x:x:x:x:x (eight 16-bit pieces) |
| | x:x:x:x:x:d.d.d.d |
| | x: [0 to FFFF]H |
| | d: [0 to 255]D |
| ipv6-prefix-length | 0 to 128 |

prefix-length — Specifies a route must match the most significant bits and have a prefix length.

Values 1 to 128

autonomous

| | |
|--------------------|--|
| Syntax | [no] autonomous |
| Context | config>router>router-advert>if>prefix |
| Description | This command specifies whether the prefix can be used for stateless address autoconfiguration. |
| Default | enabled |

on-link

| | |
|--------------------|---|
| Syntax | [no] on-link |
| Context | config>router>router-advert>if>prefix |
| Description | This command specifies whether the prefix can be used for onlink determination. |
| Default | enabled |

preferred-lifetime

| | |
|--------------------|--|
| Syntax | [no] preferred-lifetime {seconds infinite} |
| Context | config>router>router-advert>if |
| Description | This command configures the remaining length of time in seconds that this prefix will continue to be preferred, such as, time until deprecation. The address generated from a deprecated prefix should not be used as a source address in new communications, but packets received on such an interface are processed as expected. |
| Default | 604800 |
| Parameters | seconds — Specifies the remaining length of time in seconds that this prefix will continue to be preferred. infinite — Specifies that the prefix will always be preferred. A value of 4,294,967,295 represents infinity. |

valid-lifetime

| | |
|--------------------|---|
| Syntax | valid-lifetime {seconds infinite} |
| Context | config>router>router-advert>if |
| Description | This command specifies the length of time in seconds that the prefix is valid for the purpose of on-link determination. A value of all one bits (0xffffffff) represents infinity. |

The address generated from an invalidated prefix should not appear as the destination or source address of a packet.

| | |
|-------------------|--|
| Default | 2592000 |
| Parameters | <p><i>seconds</i> — Specifies the remaining length of time in seconds that this prefix will continue to be valid.</p> <p>infinite — Specifies that the prefix will always be valid. A value of 4,294,967,295 represents infinity.</p> |

reachable-time

| | |
|--------------------|---|
| Syntax | <p>reachable-time <i>milli-seconds</i></p> <p>no reachable-time</p> |
| Context | config>router>router-advert>if |
| Description | This command configures how long this router should be considered reachable by other nodes on the link after receiving a reachability confirmation. |
| Default | no reachable-time |
| Parameters | <p><i>milli-seconds</i> — Specifies the length of time the router should be considered reachable.</p> <p>Values 0 to 3600000</p> |

retransmit-time

| | |
|--------------------|--|
| Syntax | <p>retransmit-timer <i>milli-seconds</i></p> <p>no retransmit-timer</p> |
| Context | config>router>router-advert>if |
| Description | This command configures the retransmission frequency of neighbor solicitation messages. |
| Default | no retransmit-time |
| Parameters | <p><i>milli-seconds</i> — Specifies how often the retransmission should occur.</p> <p>Values 0 to 1800000</p> |

router-lifetime

| | |
|----------------|---|
| Syntax | <p>router-lifetime <i>seconds</i></p> <p>no router-lifetime</p> |
| Context | config>router>router-advert>if |

| | |
|--------------------|--|
| Description | This command sets the router lifetime. |
| Default | 1800 |
| Parameters | <i>seconds</i> — The length of time, in seconds, (relative to the time the packet is sent) that the prefix is valid for route determination. |
| Values | 0, 4 to 9000 seconds. 0 means that the router is not a default router on this link. |

use-virtual-mac

| | |
|--------------------|---|
| Syntax | [no] use-virtual-mac |
| Context | config>router>router-advert>if |
| Description | <p>This command enables sending router advertisement messages using the VRRP virtual MAC address, provided that the virtual router is currently the master.</p> <p>If the virtual router is not the master, no router advertisement messages are sent.</p> <p>The no form of the command disables sending router advertisement messages.</p> |
| Default | no use-virtual-mac |

2.13 Show, Clear, and Debug Command Reference

- [Command Hierarchies](#)
- [Command Descriptions](#)

2.13.1 Command Hierarchies

- [Show Commands](#)
- [Clear Commands](#)
- [Debug Commands](#)
- [Tools Commands](#)

2.13.1.1 Show Commands

The show L2TP commands apply only to the 7750 SR and 7450 ESS.

```
show
  — router [router-instance]
  — router service-name service-name
    — aggregate [family] [active]
    — arp [ip-int-name | ip-address/mask | mac ieee-mac-address | summary] [local |
      dynamic | static | managed]
    — authentication
      — statistics
      — statistics interface [ip-int-name | ip-address]
      — statistics policy name
    — bfd
      — bfd-template template-name
      — interface [interface-name]
      — session detail lsp-rsvp {head | tail}
      — session {ipv4 | ipv6} detail [lag lag-id] lag-port port-id
      — session lsp-name lsp-name
      — session lsp-rsvp {head | tail}
      — session {src ip-address/link-local address dest ip-address | link-local
        address} detail lsp-rsvp {head | tail} tunnel-id tunnel-id lsp-id lsp-id
      — session mpls-tp
      — session lsp-name lsp-name [link-type {cc-only | cc-cv}] detail
      — session p2mp-interface interface-name detail
      — session src ip-address/link-local address detail lsp-rsvp {head | tail} rsvp-
        session-name rsvp-session-name
      — session [src ip-address/link-local address] [ipv4 | ipv6]
      — session src ip-address/link-local address dest ip-address | link-local address
      — session src ip-address/link-local address detail
```

- **session** **summary**
- **session** **type** *type* [*ipv4* | *ipv6*]
- **dhcp**
 - **statistics** [*ip-int-name* | *ip-address*]
 - **summary**
- **dhcp6**
 - **statistics**
 - **summary**
- **ecmp**
- **fib** *slot-number* [*family*] [*ip-prefix/prefix-length* [*longer*]] [*secondary*]
- **fib** *slot-number* [*family*] **summary**
- **fib** *slot-number* *nh-table-usage*
- **fp-tunnel-table** *slot-number* [*ip-prefix/prefix-length*]
- **icmp** [*interface interface-name*]
- **icmp6** [*interface interface-name*]
- **if-attribute**
 - **srlg-group** [*name*]
- **interface** [{*ip-address* | *ip-int-name*] [**detail**] [*family*}] | **summary** | **exclude-services**
- **interface** {*ip-address* | *ip-int-name*} **eth-cfm** [**detail**]
- **interface** {*ip-address* | *ip-int-name*} **mac** [*ieee-address*]
- **interface** {*ip-address* | *ip-int-name*} **statistics**
- **interface** **dist-cpu-protection** [**detail**]
- **interface** **policy-accounting** [*class* [*index*]]
- **l2tp**
 - **eth-tunnel** [*group tunnel-group-name* [*vc-id vc-id*]]
 - **group** [*tunnel-group-name* [**statistics**]]
 - **group** **connection-id** *connection-id* [**detail**]
 - **group** [**detail**] [*session-id session-id* (*v2*)] [*state session-state*] [*peer ip-address*] [*group group-name*] [*assignment-id assignment-id*] [*local-name local-host-name*] [*remote-name remote-host-name*] [*tunnel-id tunnel-id* (*v2*)]
 - **peer** *ip-address* [**statistics**] [{*udp-port port* | *ip*}]
 - **peer** [**draining**] [{*blacklisted* | *selectable* | *unreachable*}]
 - **session** [**detail**] [*state session-state*] [*peer ip-address*] [*group group-name*] [*assignment-id assignment-id*] [*local-name local-host-name*] [*remote-name remote-host-name*] [*control-connection-id connection-id* (*v3*)]
 - **statistics**
 - **tunnel** [**statistics**] [**detail**] [*peer ip-address*] [*state tunnel-state*] [*remote-connection-id remote-connection-id* (*v3*)] [*group group-name*] [*assignment-id assignment-id*] [*local-name host-name*] [*remote-name host-name*] | **tunnel** [**statistics**] [**detail**] [*peer ip-address*] [*state tunnel-state*] [*remote-tunnel-id remote-tunnel-id* (*v2*)] [*group group-name*] [*assignment-id assignment-id*] [*local-name host-name*] [*remote-name host-name*]
 - **tunnel** *tunnel-id tunnel-id* (*v2*) [**statistics**] [**detail**]
 - **tunnel** *connection-id connection-id* (*v3*) [**statistics**] [**detail**]
- **ldp**
 - **bindings** **active**
- **mvpn**
- **neighbor** [*ip-address* | *ip-int-name* | *mac ieee-mac-address* | **summary**]
- **network-domains** [**detail**] [*network-domain-name*]
- **origin-validation**
 - **database** [*family*] [*ip-prefix/ip-prefix-length*] [*upto prefix-length2*] [*origin-as as-number*]

- **database** *[family] [ip-prefix/ip-prefix-length] [longer]*
- **database** {summary}
- **database** *[family] [static]*
- **rpki-session** *[ipv4-address] [detail]*
- **policy** *[name | damping | prefix-list name | as-path name | community name | admin]*
- **policy-edits**
- **route-table** *[family] [ip-prefix[/prefix-length]] [longer | exact | protocol protocol-name] [all]] [next-hop-type type] [qos] [alternative]*
- **route-table** *[family] summary*
- **route-table** *tunnel-endpoints [ip-prefix[/prefix-length]] [longer | exact] [detail]*
- **rtr-advertisement** *[interface interface-name] [prefix ipv6-prefix[/prefix-length] [conflicts]*
- **service-prefix**
- **sgt-qos**
 - **application** *[app-name] [dscp-dot1p]*
 - **dscp-map** *[dscp-name]*
- **static-arp** *[ip-address | ip-int-name | mac ieee-mac-addr]*
- **static-route** *[family] [[ip-prefix /mask] | [preference preference] | [next-hop ip-address] | [tag tag] [detail]*
- **status**
- **tunnel-table** *summary [ipv4 | ipv6]*
- **tunnel-table** *[protocol protocol] {ipv4 | ipv6}*
- **tunnel-table** *[ip-prefix[/mask]] [alternative] [ipv4 | ipv6] [detail]*
- **tunnel-table** *mpls-tp*
- **tunnel-table** *[ip-prefix[/mask]] protocol protocol [detail]*
- **tunnel-table** *[ip-prefix[/mask]] sdp sdp-id*
- **neighbor** *[interface-name]*

2.13.1.2 Clear Commands

- clear**
- **router** *[router-instance]*
 - **router** **service-name** *service-name*
 - **arp** {all | *ip-addr* | **interface** {*ip-int-name* | *ip-addr*}}
 - **bfd**
 - **session** *src-ip ip-address dst-ip ip-address*
 - **statistics** *src-ip ip-address dst-ip ip-address*
 - **statistics** all
 - **dhcp**
 - **statistics** *[ip-int-name | ip-address]*
 - **dhcp6**
 - **statistics** *[ip-int-name | ip-address]*
 - **forwarding-table** *[slot-number]*
 - **grt-lookup**
 - **icmp** all
 - **icmp** global
 - **icmp** **interface** *interface-name*
 - **icmp-redirect-route** {all | *ip-address*}
 - **icmp6** all
 - **icmp6** global

- **icmp6** **interface** *interface-name*
- **interface** [*ip-int-name* | *ip-address*] [**urpf-stats**] [**statistics**] [**hold-time**]
- **interface** [*ip-int-name* | *ip-address*] **policy-accounting** [**class**] [*index*]
- **interface** *ip-int-name* | *ip-address* **mac** [*ieee-address*]
- **I2tp**
 - **group** *tunnel-group-name*
 - **statistics**
 - **statistics**
 - **tunnel** *tunnel-id*
 - **statistics**
- **neighbor** {**all** | *ip-address*}
- **neighbor** [**interface** *ip-int-name* | *ip-address*]
- **router-advertisement** **all**
- **router-advertisement** [**interface** *interface-name*]
- **forwarding-table** [*slot-number*]

2.13.1.3 Debug Commands

- debug**
 - **trace**
 - **destination** *trace-destination*
 - **enable**
 - [**no**] **trace-point** [**module** *module-name*] [**type** *event-type*] [**class** *event-class*] [**task** *task-name*] [**function** *function-name*]
 - **router** [*router-instance*]
 - **router** **service-name** *service-name*
 - **ip**
 - [**no**] **arp**
 - **icmp**
 - **no icmp**
 - **icmp6** [*ip-int-name*]
 - **no icmp6**
 - [**no**] **interface** [*ip-int-name* | *ip-address*]
 - [**no**] **neighbor**
 - **packet** [*ip-int-name* | *ip-address*] [**headers**] [*protocol-id*]
 - **no packet** [*ip-int-name* | *ip-address*]
 - **route-table** [*ip-prefix/prefix-length*] [**longer**]
 - **no route-table**
 - **tunnel-table** [*ip-address*] [**ldp** | **rsvp** [*tunnel-id tunnel-id*] | **sdp** [*sdp-id sdp-id*]]
 - **I2tp**
 - **peer** *ip-address* [{**udp-port** *port* | **ip**}]
 - **mtrace**
 - [**no**] **misc**
 - [**no**] **packet** [**query** | **request** | **response**]

2.13.1.4 Tools Commands

tools

```
— dump
  — router
    — segment-routing
      — tunnel
— perform
  — router
    — l2tp
      — peer ip-address [{udp-port port | ip}]
```

2.13.2 Command Descriptions

- [Show Commands](#)
 - [L2TP Show Commands](#)
- [Clear Commands](#)
- [Debug Commands](#)
- [Tools Commands](#)

2.13.2.1 Show Commands

The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

router

| | |
|-------------|---|
| Syntax | router [<i>router-instance</i>] router service-name <i>service-name</i> |
| Context | show |
| Description | This command enables the context to display various types of information for the specified router instance. |
| Parameters | <i>router-instance</i> — specifies the router name, CPM router instance, or VPRN service ID |
| Values | |

router-instance : *router name* | *vpn-svc-id*

router-name Base | management | *cpm-vr-name* | vpls-
management

cpm-vr-name [32 characters maximum]

vpn-svc-id [1..2147483647]

Default Base

service-name — specifies the service name, up to 64 characters

Output

Sample Output: show router with PIM and S-PMSI

```
*A:Dut-D# \show router 100 pim s-pmsi
=====
PIM RSVP Spmsi tunnels
=====
P2mp  Tunnel ID  Ext Tunnel Adrs      SPMSI Index  Num          State  Multistr
e
ID                                           VPN
SGs                      am-ID
-----
100    61442      10.20.1.4             73919        8           U
p      10
=====
PIM RSVP Spmsi Interfaces : 1
=====
*A:Dut-D# \show router 100 pim s-pmsi detail
=====
PIM RSVP Spmsi tunnels
=====
P2MP ID           : 100           Tunnel ID           : 61442
Ext Tunnnel Adrs  : 10.20.1.4       Spmsi IfIndex       : 73919
Number of VPN SGs : 8           Up Time             : 0d 00:01:04
VPN Group Address : 232.100.0.0
VPN Source Address : 100.114.1.2
Up Time           : 0d 00:01:04   Multistream-Id      : 10
State              : TX Joined     Mdt Threshold        : N/A
Join Timer         : N/A           Holddown Timer       : 0d 00:00:54
VPN Group Address : 232.100.0.1
VPN Source Address : 100.114.1.2
Up Time           : 0d 00:01:04   Multistream-Id      : 10
State              : TX Joined     Mdt Threshold        : N/A
Join Timer         : N/A           Holddown Timer       : 0d 00:00:55
VPN Group Address : 232.100.0.2
VPN Source Address : 100.114.1.2
Up Time           : 0d 00:01:04   Multistream-Id      : 5
State              : TX Joined     Mdt Threshold        : N/A
Join Timer         : N/A           Holddown Timer       : 0d 00:00:53
```

aggregate

| | |
|--------------------|---|
| Syntax | aggregate [<i>family</i>] [active] |
| Context | show>router |
| Description | This command displays aggregate routes. |
| Parameters | <i>family</i> — specifies whether IPv4 or IPv6 aggregate routes are displayed Values ipv4, ipv6 active — when the active keyword is specified, inactive aggregates are filtered out |
| Output | The following output is an example of aggregate route information. |

Sample Output

```
*A:CPM133>config>router# show router aggregate
=====
Aggregates (Router: Base)
=====
Prefix                               Aggr IP-Address  Aggr AS
  Summary                            AS Set          State
  NextHop                           Community       NextHopType
-----
10.0.0.0/8                           0.0.0.0          0
  False                             False            Inactive
                                   100:33           Blackhole
-----
No. of Aggregates: 1
=====
*A:CPM133>config>router#
```

arp

| | |
|--------------------|--|
| Syntax | arp [<i>ip-int-name</i> <i>ip-address/mask</i> mac <i>ieee-mac-address</i> summary] [local dynamic static managed] |
| Context | show>router |
| Description | This command displays the router ARP table sorted by IP address. If no command line options are specified, all ARP entries are displayed. |
| Parameters | <i>ip-address/mask</i> — only displays ARP entries associated with the specified IP address and mask <i>ip-int-name</i> — only displays ARP entries associated with the specified IP interface name mac <i>ieee-mac-addr</i> — only displays ARP entries associated with the specified MAC address summary — displays an abbreviate list of ARP entries |

[**local** | **dynamic** | **static** | **managed**] — only displays ARP information associated with the keyword

Output **ARP Table Output** — The following output is an example of router ARP table information, and [Table 9](#) describes the ARP table output fields.

Sample Output

```
*B:7710-Red-RR# show router arp
=====
ARP Table (Router: Base)
=====
IP Address      MAC Address      Expiry      Type      Interface
-----
10.20.1.24      00:16:4d:23:91:b8 00h00m00s Oth      system
10.10.4.11      00:03:fa:00:d0:c9 00h57m03s Dyn[I]   to-core-sr1
10.10.4.24      00:03:fa:41:8d:20 00h00m00s Oth[I]   to-core-sr1
-----
No. of ARP Entries: 3
=====

A:ALA-A# show router ARP 10.10.0.3
=====
ARP Table
=====
IP Address      MAC Address      Expiry      Type      Interface
-----
10.10.0.3       04:5d:ff:00:00:00 00:00:00    Oth      system
=====
A:ALA-A#

A:ALA-A# show router ARP to-ser1
=====
ARP Table
=====
IP Address      MAC Address      Expiry      Type      Interface
-----
10.10.13.1      04:5b:01:01:00:02 03:53:09    Dyn      to-ser1
=====
A:ALA-A#
```

Table 9 **ARP Fields**

| Label | Description |
|-------------|-----------------------------------|
| IP Address | The IP address of the ARP entry. |
| MAC Address | The MAC address of the ARP entry. |
| Expiry | The age of the ARP entry. |

Table 9 ARP Fields (Continued)

| Label | Description |
|--------------------|---|
| Type | <p>Dyn The ARP entry is a dynamic ARP entry.</p> <p>Inv The ARP entry is an inactive static ARP entry (invalid).</p> <p>Oth The ARP entry is a local or system ARP entry.</p> <p>Sta The ARP entry is an active static ARP entry.</p> |
| *Man | The ARP entry is a managed ARP entry. |
| Int | The ARP entry is an internal ARP entry. |
| [I] | The ARP entry is in use. |
| Interface | The IP interface name associated with the ARP entry. |
| No. of ARP Entries | The number of ARP entries displayed in the list. |

authentication

| | |
|--------------------|--|
| Syntax | authentication |
| Context | show>router |
| Description | This command enables the command to display authentication statistics. |

statistics

| | |
|--------------------|--|
| Syntax | statistics statistics interface [<i>ip-int-name</i> <i>ip-address</i>] statistics policy <i>name</i> |
| Context | show>router>authentication |
| Description | This command displays interface or policy authentication statistics. |
| Parameters | interface [<i>ip-int-name</i> <i>ip-address</i>] — specifies an existing interface name or IP address Values <i>ip-int-name</i> : 32 chars max <i>ip-address</i> : a.b.c.d policy name — specifies an existing policy name |

Output **Authentication Statistics Output** — The following output is an example of authentication statistics, and [Table 10](#) describes the fields.

Sample Output

```
A:ALU-3>show>router>auth# statistics
=====
Authentication Global Statistics
=====
Client Packets Authenticate Fail      : 0
Client Packets Authenticate Ok       : 12
=====
```

Table 10 **Authentication Statistics Fields**

| Label | Description |
|----------------------------------|---|
| Client Packets Authenticate Fail | The number of packets that failed authentication. |
| Client Packets Authenticate Ok | The number of packets that were authenticated. |

bfd

Syntax **bfd**

Context show>router

Description This command enables the context to display bidirectional forwarding detection (BFD) information.

Output The following output is an example of BFD information.

Sample Output

```
*A:Dut-D# show router 3 bfd session
=====
BFD Session
=====
InterfaceState      Tx Intvl  Rx Intvl  Multipl
Remote Address      Protocols      Tx Pkts   Rx Pkts   Type
-----
ies-3-121.1.3.3      Up (3)                10        10        3
121.1.3.2            ospf2            N/A        N/A        cpm-np
ies-3-122.1.4.3      Up (3)                100       100        3
122.1.4.2            pim              455       464        iom
-----
No. of BFD sessions: 2
=====
*A:Dut-D#
```



```
*A:Dut-C# show router bfd session src 11.120.1.4 dest 11.120.1.3
=====
BFD Session
=====
Remote Address : 11.120.1.3
Admin State    : Up                               Oper State     : Up (3)
Protocols      : static
Rx Interval    : 10                               Tx Interval    : 10
Multiplier    : 3                               Echo Interval  : 0
Up Time        : 1d 19:03:28                     Up Transitions : 2
Down Time      : None                             Down Transitions : 1
                                                    Version Mismatch : 0

Forwarding Information
Local Discr    : 19269                           Local State    : Up (3)
Local Diag     : 0 (None)                         Local Mode     : Async
Local Min Tx   : 10                               Local Mult     : 3
Last Sent (ms) : 6                               Local Min Rx   : 10
Type           : cpm-np
Remote Discr   : 5101                             Remote State   : Up (3)
Remote Diag    : 0 (None)                         Remote Mode    : Async
Remote Min Tx  : 1000                             Remote Mult    : 3
Last Recv (ms) : 367                             Remote Min Rx  : 10
=====
*A:Dut-C#
```

bfd-template

| | |
|--------------------|---|
| Syntax | bfd-template <i>template-name</i> |
| Context | show>router>bfd |
| Description | This command displays BFD template information. |
| Output | The following output is an example of BFD template information. |

Sample Output

```
*A:mlstp-dutA# show router bfd bfd-template "privatebed-bfd-template"
=====
BFD Template privatebed-bfd-template
=====
Template Name      : privatebed-* Template Type      : cpmNp
Transmit Timer     : 10 msec      Receive Timer      : 10 msec
CV Transmit Interval : 1000 msec
Template Multiplier : 3           Echo Receive Interval : 100 msec

Mpls-tp Association
privatebed-oam-template
=====
* indicates that the corresponding row element may have been truncated.
*A:mlstp-dutA# show router bfd session
```

```

=====
BFD Session
=====
Interface/Lsp Name      State      Tx Intvl  Rx Intvl  Multipl
  Remote Address/Info    Protocols      Tx Pkts   Rx Pkts   Type
-----
wp::lsp-32              Down (1)      1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
wp::lsp-33              Down (1)      1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
wp::lsp-34              Down (1)      1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
wp::lsp-35              Down (1)      1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
wp::lsp-36              Down (1)      1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
wp::lsp-37              Down (1)      1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
wp::lsp-38              Down (1)      1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
wp::lsp-39              Down (1)      1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
wp::lsp-40              Down (1)      1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
wp::lsp-41              Down (1)      1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
pp::lsp-32              Up (3)        1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
pp::lsp-33              Up (3)        1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
pp::lsp-34              Up (3)        1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
pp::lsp-35              Up (3)        1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
pp::lsp-36              Up (3)        1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
pp::lsp-37              Up (3)        1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
pp::lsp-38              Up (3)        1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
pp::lsp-39              Up (3)        1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
pp::lsp-40              Up (3)        1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
pp::lsp-41              Up (3)        1000      1000      3
  0::0.0.0.0            mplsTp        N/A       N/A       cpm-np
-----
No. of BFD sessions: 20
-----
wp = Working path   pp = Protecting path
=====

```

interface

| | |
|--------------------|---|
| Syntax | interface <i>[interface-name]</i> |
| Context | show>router>bfd |
| Description | This command displays interface information. |
| Output | The following output is an example of BFD interface information, and Table 11 describes the fields. |

Sample Output

```
*A:Dut-B# show router bfd interface
=====
BFD Interface
=====
Interface name          Tx Interval    Rx Interval    Multiplier
-----
port-1-1                500            500            3
port-1-1                10             10             3
port-1-2                500            500            3
port-1-2                10             10             3
port-1-3                500            500            3
port-1-3                10             10             3
port-1-4                500            500            3
port-1-4                10             10             3
port-1-5                500            500            3
...
=====
```

Table 11 BFD Interface Fields

| Label | Description |
|-------------|---|
| TX Interval | Displays the interval, in milliseconds, between the transmitted BFD messages to maintain the session. |
| RX Interval | Displays the expected interval, in milliseconds, between the received BFD messages to maintain the session. |
| Multiplier | Displays the integer used by BFD to declare when the neighbor is down. |

session

| | |
|---------------|---|
| Syntax | session detail <i>lsp-rsvp {head tail}</i> session { <i>ipv4 ipv6</i> } detail [<i>lag lag-id</i>] lag-port <i>port-id</i> session <i>lsp-name</i> <i>Lsp Name</i> session lsp-rsvp { <i>head tail</i> } |
|---------------|---|

```

session src ip-address/link-local address dest ip-address | link-local address detail lsp-rsvp {head | tail} tunnel-id tunnel-id lsp-id lsp-id
session mpls-tp
session lsp-name Lsp Name [link-type {cc-only | cc-cv}] detail
session p2mp-interface interface-name detail
session src ip-address/link-local address detail lsp-rsvp {head | tail} rsvp-session-name rsvp-session-name
session [src ip-address/link-local address] [ipv4 | ipv6]
session src ip-address/link-local address dest ip-address | link-local address
session src ip-address/link-local address detail
session summary
session type type [ipv4 | ipv6]

```

Context show>router>bfd

Description This command displays session information.

Parameters *ip-address* — only displays the interface information associated with the specified IP address

Values ipv4-address: a.b.c.d (host bits must be 0)

type — specifies the session type

Values iom | central | cpm-np

Output The following output is an example of BFD session information, and [Table 12](#) describes the fields.

Sample Output

```

A:Dut-B# show router bfd session
=====
BFD Session
=====
Interface          State      Tx Intvl  Rx Intvl  Multipl
Remote Address     Protocols  Tx Pkts   Rx Pkts   Type
-----
port-1-1           Up (3)     500        500        3
10.1.1.3           pim isis   50971      50718      iom
port-1-1           Up (3)     10         10         3
3FFE::A01:103      static bgp  N/A        N/A        cpm-np
port-1-1           Up (3)     10         10         3
FE80::A0A:A03      pim isis   N/A        N/A        cpm-np
port-1-2           Up (3)     500        500        3
10.2.1.3           pim isis   50968      50718      iom
port-1-2           Up (3)     10         10         3
3FFE::A02:103      static bgp  N/A        N/A        cpm-np
port-1-2           Up (3)     10         10         3
...
=====
*A:Dut-B#

```

```
A:Dut-B# show router bfd session src 3FFE::A01:102 dest 3FFE::A01:103
=====
BFD Session
=====
Remote Address : 3FFE::A01:103
Admin State    : Up                               Oper State     : Up (3)
Protocols      : static bgp
Rx Interval    : 10                               Tx Interval    : 10
Multiplier     : 3                               Echo Interval  : 0
Up Time        : 0d 07:24:54                     Up Transitions : 1
Down Time      : None                             Down Transitions : 0
                                           Version Mismatch : 0

Forwarding Information
Local Discr    : 2051                             Local State    : Up (3)
Local Diag     : 0 (None)                         Local Mode     : Async
Local Min Tx   : 10                               Local Mult     : 3
Last Sent (ms) : 5                               Local Min Rx   : 10
Type           : cpm-np
Remote Discr   : 1885                             Remote State   : Up (3)
Remote Diag    : 0 (None)                         Remote Mode    : Async
Remote Min Tx  : 10                               Remote Mult    : 3
Last Recv (ms) : 1                               Remote Min Rx  : 10
=====
A:Dut-B#

*A:Dut-B# show router bfd session src FE80::A0A:A02-port-1-10 dest FE80::A0A:A03-
port-1-10
=====
BFD Session
=====
Remote Address : FE80::A0A:A03
Admin State    : Up                               Oper State     : Up (3)
Protocols      : pim isis ospf3
Rx Interval    : 10                               Tx Interval    : 10
Multiplier     : 3                               Echo Interval  : 0
Up Time        : 0d 07:10:20                     Up Transitions : 3
Down Time      : None                             Down Transitions : 2
                                           Version Mismatch : 0

Forwarding Information
Local Discr    : 42                               Local State    : Up (3)
Local Diag     : 3 (Neighbor signalled s* Local Mode     : Async
Local Min Tx   : 10                               Local Mult     : 3
Last Sent (ms) : 6                               Local Min Rx   : 10
Type           : cpm-np
Remote Discr   : 270                             Remote State   : Up (3)
Remote Diag    : 0 (None)                         Remote Mode    : Async
Remote Min Tx  : 10                               Remote Mult    : 3
Last Recv (ms) : 8                               Remote Min Rx  : 10
=====
* indicates that the corresponding row element may have been truncated.
A:Dut-D#

*A:Dut-B# show router bfd session ipv4
=====
BFD Session
=====
Interface                               State          Tx Intvl  Rx Intvl  Multipl
```

```

Remote Address      Protocols      Tx Pkts  Rx Pkts  Type
-----
port-1-1            Up (3)         500      500      3
  10.1.1.3          pim isis       51532    51279    iom
port-1-2            Up (3)         500      500      3
  10.2.1.3          pim isis       51529    51279    iom
port-1-3            Up (3)         500      500      3
  10.3.1.3          pim isis       51529    51279    iom
port-1-4            Up (3)         500      500      3
  10.4.1.3          pim isis       51529    51279    iom
port-1-5            Up (3)         500      500      3
  10.5.1.3          pim isis       51529    51279    iom
port-1-6            Up (3)         500      500      3
  10.6.1.3          pim isis       51529    51279    iom
...
=====
*A:Dut-B#

```

```

*A:Dut-B# show router bfd session ipv6
=====
BFD Session
=====
Interface      State      Tx Intvl  Rx Intvl  Multipl
Remote Address Protocols  Tx Pkts   Rx Pkts   Type
-----
port-1-1        Up (3)      10         10         3
  3FFE::A01:103 static bgp   N/A        N/A        cpm-np
port-1-1        Up (3)      10         10         3
  FE80::A0A:A03 pim isis    N/A        N/A        cpm-np
port-1-2        Up (3)      10         10         3
  3FFE::A02:103 static bgp   N/A        N/A        cpm-np
port-1-2        Up (3)      10         10         3
  FE80::A0A:A03 pim isis    N/A        N/A        cpm-np
port-1-3        Up (3)      10         10         3
  3FFE::A03:103 static bgp   N/A        N/A        cpm-np
port-1-3        Up (3)      10         10         3
  FE80::A0A:A03 pim isis    N/A        N/A        cpm-np
port-1-4        Up (3)      10         10         3
  3FFE::A04:103 static bgp   N/A        N/A        cpm-np
port-1-4        Up (3)      10         10         3
...
=====
*A:Dut-B#

```

```

*A:Dut-D# show router bfd session summary
=====
BFD Session Summary
=====
Termination    Session Count
-----
central        0
cpm-np         500
iom, slot 1    0
iom, slot 2    0
iom, slot 3    250
iom, slot 4    0
iom, slot 5    0

```

```

Total                               750
=====
*A:Dut-D#

*A:Dut-B# show router bfd session detail lsp-
rsvp head src 10.20.1.2 dest 10.20.1.5 tunnel-id 1 lsp-id 31744
=====
BFD On LSP Session
=====
Rsvp Session Name : lsp1::path1
Remote Address : 10.20.1.5
Lsp Id          : 31744          Tunnel Id          : 1
Oper State      : Up             Protocols         : lsp
Recd Msgs       : 240            Sent Msgs         : 240
Up Time         : 0d 00:03:58    Up Transitions    : 1
Down Time       : None           Down Transitions  : 0
                                   Version Mismatch     : 0

Forwarding Information

Local Discr      : 1              Local State       : Up
Local Diag       : 0 (None)
Local Mode       : Async
Local Min Tx     : 1000           Local Mult        : 3
Last Sent        : 07/28/2015 19:05:13 Local Min Rx      : 1000
Type             : central
Remote Discr     : 1              Remote State      : Up
Remote Diag      : 0 (None)       Remote Mode       : Async
Remote Min Tx    : 1000           Remote Mult       : 3
Last Recv        : 07/28/2015 19:05:13 Remote Min Rx    : 1000
=====
=====

```

Table 12 BFD Session Field Descriptions

| Label | Description |
|----------|---|
| State | Displays the administrative state for this BFD session. |
| Protocol | Displays the active protocol. |
| Tx Intvl | Displays the interval, in milliseconds, between the transmitted BFD messages to maintain the session. |
| Tx Pkts | Displays the number of transmitted BFD packets. |
| Rx Intvl | Displays the expected interval, in milliseconds, between the received BFD messages to maintain the session. |
| Rx Pkts | Displays the number of received packets. |
| Mult | Displays the integer used by BFD to declare when the neighbor is down. |

dhcp

| | |
|--------------------|---|
| Syntax | dhcp |
| Context | show>router |
| Description | This command enables the context to display DHCP related information. |

dhcp6

| | |
|--------------------|--|
| Syntax | dhcp6 |
| Context | show>router |
| Description | This command enables the context to display DHCP6 related information. |

statistics

| | |
|--------------------|--|
| Syntax | statistics [interface <i>ip-int-name</i> <i>ip-address</i>] |
| Context | show>router>dhcp |
| Description | <p>This command displays statistics for DHCP relay and DHCP snooping.</p> <p>If no IP address or interface name is specified, then all configured interfaces are displayed.</p> <p>If an IP address or interface name is specified, then only data regarding the specified interface is displayed.</p> |
| Parameters | <i>ip-int-name</i> <i>ip-address</i> — displays statistics for the specified IP interface |
| Output | The following output is an example of DHCP statistics information, and Table 13 describes the output fields. |

Sample Output

```
A:ALA-1# show router dhcp statistics
=====
DHCP statistics (Router: Base)
=====
Msg-type           Rx           Tx           Dropped
-----
1 SOLICIT           0            0            0
2 ADVERTISE         0            0            0
3 REQUEST           0            0            0
4 CONFIRM           0            0            0
5 RENEW             0            0            0
6 REBIND            0            0            0
7 REPLY             0            0            0
8 RELEASE           0            0            0
9 DECLINE           0            0            0
```



```

10 RECONFIGURE                0                0                0
11 INFO_REQUEST                0                0                0
12 RELAY_FORW                  0                0                0
13 RELAY_REPLY                  0                0                0
-----
Dhcp Drop Reason Counters :
-----
 1 Dhcp6 oper state is not Up on src itf                0
 2 Dhcp6 oper state is not Up on dst itf                0
 3 Relay Reply Msg on Client Itf                        0
 4 Hop Count Limit reached                             0
 5 Missing Relay Msg option, or illegal msg type        0
 6 Unable to determine destination client Itf           0
 7 Out of Memory                                        0
 8 No global Pfx on Client Itf                          0
 9 Unable to determine src Ip Addr                      0
10 No route to server                                  0
11 Subscr. Mgmt. Update failed                          0
12 Received Relay Forw Message                         0
13 Packet too small to contain valid dhcp6 msg          0
14 Server cannot respond to this message               0
15 No Server Id option in msg from server               0
16 Missing or illegal Client Id option in client msg    0
17 Server Id option in client msg                      0
18 Server DUID in client msg does not match our own     0
19 Client sent message to unicast while not allowed     0
20 Client sent message with illegal src Ip address      0
21 Client message type not supported in pfx delegation  0
22 Nbr of addrs or pfxs exceeds allowed max (128) in msg 0
23 Unable to resolve client's mac address               0
24 The Client was assigned an illegal address           0
25 Illegal msg encoding                                0
=====
A:ALA-1#

```

Table 13 DHCP Statistics Fields

| Label | Description |
|----------------------------|---|
| Received Packets | The number of packets received from the DHCP clients. |
| Transmitted Packets | The number of packets transmitted to the DHCP clients. |
| Received Malformed Packets | The number of malformed packets received from the DHCP clients. |
| Received Untrusted Packets | The number of untrusted packets received from the DHCP clients. |
| Client Packets Discarded | The number of packets received from the DHCP clients that were discarded. |
| Client Packets Relayed | The number of packets received from the DHCP clients that were forwarded. |

Table 13 DHCP Statistics Fields (Continued)

| Label | Description |
|--------------------------|--|
| Client Packets Snooped | The number of packets received from the DHCP clients that were snooped. |
| Server Packets Discarded | The number of packets received from the DHCP server that were discarded. |
| Server Packets Relayed | The number of packets received from the DHCP server that were forwarded. |
| Server Packets Snooped | The number of packets received from the DHCP server that were snooped. |

statistics

Syntax **statistics**

Context show>router>dhcp6

Description This command displays statistics for DHCP relay and DHCP snooping.

Output The following output is an example of DHCP statistics information.

Sample Output

```
A:ALA-1# show router dhcp6 statistics
=====
DHCP6 statistics (Router: Base)
=====
Msg-type           Rx           Tx           Dropped
-----
1 SOLICIT           0           0           0
2 ADVERTISE         0           0           0
3 REQUEST           0           0           0
4 CONFIRM           0           0           0
5 RENEW             0           0           0
6 REBIND            0           0           0
7 REPLY             0           0           0
8 RELEASE           0           0           0
9 DECLINE           0           0           0
10 RECONFIGURE      0           0           0
11 INFO_REQUEST     0           0           0
12 RELAY_FORW       0           0           0
13 RELAY_REPLY      0           0           0
-----
Dhcp6 Drop Reason Counters :
-----
1 Dhcp6 oper state is not Up on src itf      0
2 Dhcp6 oper state is not Up on dst itf      0
3 Relay Reply Msg on Client Itf              0
4 Hop Count Limit reached                    0
```

```

5 Missing Relay Msg option, or illegal msg type 0
6 Unable to determine destination client Itf 0
7 Out of Memory 0
8 No global Pfx on Client Itf 0
9 Unable to determine src Ip Addr 0
10 No route to server 0
11 Subscr. Mgmt. Update failed 0
12 Received Relay Forw Message 0
13 Packet too small to contain valid dhcp6 msg 0
14 Server cannot respond to this message 0
15 No Server Id option in msg from server 0
16 Missing or illegal Client Id option in client msg 0
17 Server Id option in client msg 0
18 Server DUID in client msg does not match our own 0
19 Client sent message to unicast while not allowed 0
20 Client sent message with illegal src Ip address 0
21 Client message type not supported in pfx delegation 0
22 Nbr of addrs or pfxs exceeds allowed max (128) in msg 0
23 Unable to resolve client's mac address 0
24 The Client was assigned an illegal address 0
25 Illegal msg encoding 0
=====
A:ALA-1#

```

summary

| | |
|--------------------|--|
| Syntax | summary |
| Context | show>router>dhcp |
| Description | Display the status of the DHCP Relay and DHCP Snooping functions on each interface. |
| Output | The following output is an example of DHCP summary information, and Table 14 describes the output fields for DHCP summary. |

Sample Output

```

A:ALA-1# show router dhcp summary
=====
DHCP6 Summary (Router: Base)
=====
Interface Name      Nbr      Used/Max Relay  Admin  Oper Relay
  SapId      Resol.  Used/Max Server Admin  Oper Server
-----
interfaceServiceDefault  No      0/0      Up      NoServerCo*
  sap:1/2/12:1      0/8000      Up      Up
interfaceService      No      0/0      Down      Down
  sap:1/2/1      0/8000      Down      Down
interfaceServiceNonDefault  No      0/0      Up      NoServerCo*
  sap:1/2/12:2      0/8000      Down      Down
ip-61.4.113.4      Yes      575/8000      Up      Up
  sap:1/1/1:1      580/8000      Up      Up
=====
A:ALA-1#

```

Table 14 DHCP Summary Field Descriptions

| Label | Description |
|----------------|--|
| Interface Name | Name of the router interface. |
| Info Option | Indicates whether Option 82 processing is enabled on the interface. |
| Auto Filter | Indicates whether IP Auto Filter is enabled on the interface. |
| Snoop | Indicates whether Auto ARP table population is enabled on the interface. |
| Interfaces | Indicates the total number of router interfaces on the router. |

ecmp

- Syntax** `ecmp`
- Context** `show>router`
- Description** This command displays the ECMP settings for the router.
- Output** The following output is an example of ECMP settings information, and [Table 15](#) describes the output fields for the router ECMP settings.

Sample Output

```

A:ALA-A# show router ecmp
=====
Router ECMP
=====
Instance      Router Name      ECMP      Configured-ECMP-Routes
-----
1             Base             True      8
=====
A:ALA-A#
*A:Dut-C# show router ecmp

=====
Router ECMP
=====
Instance      Router Name      ECMP      Max-ECMP-      Weight ECMP
              Router Name      Rtes
-----
1             Base             True      32             True
=====

```

Table 15 ECMP Fields

| Label | Description |
|------------------------|--|
| Instance | The router instance number. |
| Router Name | The name of the router instance. |
| ECMP | False ECMP is disabled for the instance. |
| | True ECMP is enabled for the instance. |
| Configured-ECMP-Routes | The number of ECMP routes configured for path sharing. |

fib

Syntax **fib** *slot-number* [*family*] [*ip-prefix/prefix-length*] [**longer**] [**secondary**] [**exclude-services**]
fib *slot-number* [*family*] **summary**
fib *slot-number* **nh-table-usage**

Context show>router

Description This command displays the active FIB entries for a specific IOM or linecard.

Parameters *slot-number* — displays routes only matching the specified chassis slot number

Default all IOMs

Values 1 to 10

family — displays the router IP interface table to display

Values **ipv4** — displays only those peers that have the IPv4 family enabled

ipv6 — displays the peers that are IPv6-capable

ip-prefix/prefix-length — displays FIB entries only matching the specified *ip-prefix* and length

Values The following values apply to the 7450 ESS:

ipv4-prefix: a.b.c.d (host bits must be 0)

ipv4-prefix-length: 0 to 32

Values The following values apply to the 7750 SR and 7950 XRS:

ipv4-prefix: a.b.c.d (host bits must be 0)

ipv4-prefix-length: 0 to 32

| | |
|---------------------|---|
| ipv6-prefix: | x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D |
| ipv6-prefix-length: | 0 to 128 |

- longer** — displays FIB entries matching the *ip-prefix/mask* and routes with longer masks
- secondary** — displays secondary VRF ID information
- summary** — displays summary FIB information for the specified slot number
- nh-table-usage** — displays next-hop table usage

Output The following output is an example of FIB information.

Sample Output

```
show router fib 1 131.132.133.134/32
=====
FIB Display
=====
Prefix                                     Protocol
  NextHop
-----
131.132.133.134/32                        OSPF
    66.66.66.66 (loop7)
    Next-hop type: tunneled, Owner: RSVP, Tunnel-ID: <out-ifindex-from-route>
-----
Total Entries : 1
=====

*A:Dut-C# show router fib 1 1.1.1.1/32
=====
FIB Display
=====
Prefix                                     Protocol
  NextHop
-----
1.1.1.1/32                                BGP
    10.20.1.1 (Transport:RSVP LSP:1)
-----
Total Entries : 1
=====

*A:Dut-C# show router fib 1
=====
FIB Display
=====
Prefix                                     Protocol
  NextHop
-----
1.1.2.0/24                                ISIS
    1.1.3.1 (to_Dut-A)
```

```

        1.2.3.2 (to_Dut-B)
1.1.3.0/24                                LOCAL
        1.1.3.0 (to_Dut-A)
1.1.9.0/24                                ISIS
        1.1.3.1 (to_Dut-A)
1.2.3.0/24                                LOCAL
        1.2.3.0 (to_Dut-B)
1.2.9.0/24                                ISIS
        1.2.3.2 (to_Dut-B)
10.12.0.0/24                              LOCAL
        10.12.0.0 (itfToArborCP_02)
10.20.1.1/32                              ISIS
        1.1.3.1 (to_Dut-A)
10.20.1.2/32                              ISIS
        1.2.3.2 (to_Dut-B)
10.20.1.3/32                              LOCAL
        10.20.1.3 (system)
20.12.0.43/32                             STATIC
        vprn1:mda-1-1
20.12.0.44/32                             STATIC
        vprn1:mda-2-1
20.12.0.45/32                             STATIC
        vprn1:mda-2-2
20.12.0.46/32                             STATIC
        vprn1:mda-3-1
138.203.71.202/32                         STATIC
        10.12.0.2 (itfToArborCP_02)
-----
Total Entries : 15
-----
=====

*A:Dut-C>config>router>mpls>lsp# show router fib 1 5.3.0.1/32 extensive
=====
FIB Display (Router: Base)
=====
Dest Prefix          : 5.3.0.1/32
Protocol             : BGP
Indirect Next-Hop    : 10.0.0.1
  QoS                : Priority=n/c, FC=n/c
  Source-Class       : 0
  Dest-Class         : 0
  ECMP-Weight        : 1
  Resolving Next-Hop : 1.0.0.2 (RSVP tunnel:115)
    ECMP-Weight      : 1
  Resolving Next-Hop : 1.0.0.2 (RSVP tunnel:61443)
    ECMP-Weight      : 1
Indirect Next-Hop    : 10.0.0.2
  QoS                : Priority=n/c, FC=n/c
  Source-Class       : 0
  Dest-Class         : 0
  ECMP-Weight        : 30
  Resolving Next-Hop : 1.0.0.3 (RSVP tunnel:94)
    ECMP-Weight      : 20
  Resolving Next-Hop : 1.0.0.3 (RSVP tunnel:61442)
    ECMP-Weight      : 1
=====
Total Entries : 1
=====

```

```
*A:Dut-C> show router fib 1 10.0.0.2/32 extensive
=====
FIB Display (Router: Base)
=====
Dest Prefix          : 10.0.0.2/32
Protocol             : OSPF
Next-Hop             : 1.0.0.3 (RSVP tunnel:94)
  QoS                : Priority=n/c, FC=n/c
  Source-Class       : 0
  Dest-Class         : 0
  ECMP-Weight        : 20
Next-Hop             : 1.0.0.3 (RSVP tunnel:61442)
  QoS                : Priority=n/c, FC=n/c
  Source-Class       : 0
  Dest-Class         : 0
  ECMP-Weight        : 1
=====
Total Entries : 1
=====

*A:Dut-C> show router route-table 10.1.0.5/32 extensive
=====
Route Table (Router: Base)
=====
Dest Prefix          : 10.1.0.5/32
Protocol             : STATIC
Age                 : 00h01m37s
Preference          : 5
Next-Hop            : 1.0.0.2 (RSVP tunnel:128)
  QoS               : Priority=n/c, FC=n/c
  Source-Class      : 0
  Dest-Class        : 0
  Metric            : 1
  ECMP-Weight       : 10
Next-Hop            : 1.0.0.2 (RSVP tunnel:132)
  QoS               : Priority=n/c, FC=n/c
  Source-Class      : 0
  Dest-Class        : 0
  Metric            : 1
  ECMP-Weight       : 1
-----
No. of Destinations: 1
=====

*A:Dut-C> show router fib 1 10.1.0.5/32 extensive
=====
FIB Display (Router: Base)
=====
Dest Prefix          : 10.1.0.5/32
Protocol             : STATIC
Next-Hop            : 1.0.0.2 (RSVP tunnel:128)
  QoS               : Priority=n/c, FC=n/c
  Source-Class      : 0
  Dest-Class        : 0
  ECMP-Weight       : 10
Next-Hop            : 1.0.0.2 (RSVP tunnel:132)
  QoS               : Priority=n/c, FC=n/c
  Source-Class      : 0
```



```

      Dest-Class          : 0
      ECMP-Weight         : 1
=====
Total Entries : 1
=====

*A:Dut-B# show router fib 1 10.15.1.0/24
=====
FIB Display
=====
Prefix [Flags]                                Protocol
  NextHop
-----
10.15.1.0/24                                BGP
   10.20.1.3 (Transport:SR)
-----
Total Entries : 1
=====

*A:Dut-B# show router fib 1 10.15.1.0/24 extensive
=====
FIB Display (Router: Base)
=====
Dest Prefix          : 10.15.1.0/24
Protocol             : BGP
Installed            : Y
Indirect Next-Hop    : 10.20.1.3
  Label              : 262123
  QoS                 : Priority=n/c, FC=n/c
  Source-Class        : 0
  Dest-Class          : 0
  ECMP-Weight         : 1
  Resolving Next-Hop  : 10.20.1.3 (SR tunnel)
    ECMP-Weight       : 1
=====
Total Entries : 1
=====
```

fp-tunnel-table

| | |
|--------------------|---|
| Syntax | fp-tunnel-table <i>slot-number</i> [<i>ip-prefix/prefix-length</i>] |
| Context | show>router |
| Description | <p>This command displays the IOM/IMM label, next-hop and outgoing interface information for BGP, LDP and RSVP tunnels used in any of the following applications:</p> <ul style="list-style-type: none">• BGP shortcut (config>router>bgp>next-hop-resolution>shortcut-tunnel)• IGP shortcut (config>router>isis[ospf]>igp-shortcut)• IGP prefix resolved to an LDP LSP (config>router>ldp-shortcut)• Static route resolved to a LDP or RSVP LSP• VPRN auto-bind |

- 6PE/6VPE.

| | | | |
|-------------------|---|--------------------------------------|-------------------------------------|
| Parameters | <i>slot-number</i> — displays information for the specified slot | | |
| | Values | 1 to 10 | |
| | <i>ip-prefix[/prefix-length]</i> — displays routes only matching the specified <i>ip-address</i> and length | | |
| | Values | | |
| | ipv4-prefix: | a.b.c.d (host bits must be set to 0) | |
| | ipv4-prefix-length: | 0 to 32 | |
| | ipv6 | ipv6-prefix[/pref*: | x:x:x:x:x:x:x (eight 16-bit pieces) |
| | | | x:x:x:x:x:x:d.d.d.d |
| | | | x: [0 to FFFF]H |
| | | | d: [0 to 255]D |
| | | prefix-length: | 1 to 128 |

Output The following output is an example of router FP tunnel information.

Sample Output

```
*A:Dut-B# show router fp-tunnel-table 1 10.20.1.3/32
=====
Tunnel Table Display

Legend:
B - FRR Backup
=====
Destination                                Protocol  Tunnel-ID
      Lbl                                NextHop   Intf/Tunnel
-----
10.20.1.3/32                               LDP      -
      262137                             10.2.1.3   1/1/3:1
10.20.1.3/32                               RSVP      1
      262133                             10.2.1.3   1/1/3:1
10.20.1.3/32                               SR-ISIS-0 -
      18602                              10.2.1.3   1/1/3:1
10.20.1.3/32                               SR-OSPF-0 -
      19102                              10.2.1.3   1/1/3:1
-----
Total Entries : 4
=====
*A:Dut-B#

*A:Dut-C# show router fp-tunnel-table 1
=====
Tunnel Table Display

Legend:
B - FRR Backup
```

```
=====
Destination                               Protocol  Tunnel-ID
      Lbl                               NextHop   Intf/Tunnel
-----
4.0.0.1/32                               SR-ISIS-0  -
      20001                             1.3.4.4    2/1/3:1
      20001/21005                         1.2.3.2(B) 1/1/2
10.20.1.2/32                             SR-ISIS-0  -
      21002                             1.2.3.2    1/1/2
      21002/21005                         1.3.4.4(B) 2/1/3:1
10.20.1.4/32                             SR-ISIS-0  -
      21004                             1.3.4.4    2/1/3:1
      21004/21005                         1.2.3.2(B) 1/1/2
10.20.1.5/32                             SR-ISIS-0  -
      21005                             1.2.3.2    1/1/2
      21005                             1.3.4.4(B) 2/1/3:1
-----
Total Entries : 4
-----
=====
*A:Dut-C#

*A:Dut-C# show router fp-tunnel-table 1
=====
Tunnel Table Display

Legend:
B - FRR Backup
=====
Destination                               Protocol  Tunnel-ID
      Lbl                               NextHop   Intf/Tunnel
-----
1.1.3.1/32                               SR        -
      3                                 1.1.3.1    1/1/1
1.2.3.2/32                               SR        -
      3                                 1.2.3.2    1/1/2:1
1.3.5.5/32                               SR        -
      3                                 1.3.5.5    2/1/1
2.2.3.2/32                               SR        -
      3                                 2.2.3.2    1/1/2:2
10.20.1.1/32                             SR-OSPF-0  -
      21011                             1.1.3.1    1/1/1
      22011                             1.2.3.2(B) 1/1/2:1
10.20.1.2/32                             SR-OSPF-0  -
      22022                             2.2.3.2    1/1/2:2
      24022/25044                       1.3.5.5(B) 2/1/1
10.20.1.4/32                             SR-OSPF-0  -
      25044                             1.3.5.5    2/1/1
      22044                             2.2.3.2    1/1/2:2
10.20.1.5/32                             SR-OSPF-0  -
      25055                             1.3.5.5    2/1/1
      24055/22044                       2.2.3.2(B) 1/1/2:2
10.20.1.6/32                             SR-OSPF-0  -
      25066                             1.3.5.5    2/1/1
      24066/22044                       2.2.3.2(B) 1/1/2:2
-----
Total Entries : 9
-----
=====
```

*A:Dut-C#

*A:Dut-F# show router fp-tunnel-table 1

=====

Tunnel Table Display

Legend:

B - FRR Backup

=====

| Destination Lbl | NextHop | Intf/Tunnel | Protocol | Tunnel-ID |
|--------------------|--------------|-------------|-----------|-----------|
| ----- | | | | |
| 1.0.11.1/32 | | | SR-OSPF-0 | - |
| 30004 | 1.0.26.2 | 1/1/3:1 | | |
| 40004 | 1.0.36.3 (B) | 1/1/4:1 | | |
| 1.0.22.2/32 | | | SR-OSPF-0 | - |
| 30005 | 1.0.26.2 | 1/1/3:1 | | |
| 20005/40004 | 1.0.36.3 (B) | 1/1/4:1 | | |
| 1.0.26.2/32 | | | SR | - |
| 3 | 1.0.26.2 | 1/1/3:1 | | |
| 50011/60001 | 1.0.56.5 (B) | 1/1/2:1 | | |
| 1.0.26.2/32 | | | SR | - |
| 3 | 1.0.26.2 | 1/1/3:1 | | |
| 20005/40004 | 1.0.36.3 (B) | 1/1/4:1 | | |
| 1.0.33.3/32 | | | SR-OSPF-0 | - |
| 40000 | 1.0.36.3 | 1/1/4:1 | | |
| 30998 | 1.0.26.2 (B) | 1/1/3:1 | | |
| 1.0.36.3/32 | | | SR | - |
| 3 | 1.0.36.3 | 1/1/4:1 | | |
| 1.0.44.4/32 | | | SR-OSPF-0 | - |
| 30001 | 1.0.26.2 | 1/1/3:1 | | |
| 60001 | 1.0.56.5 (B) | 1/1/2:1 | | |
| 1.0.55.5/32 | | | SR-OSPF-0 | - |
| 60002 | 1.0.56.5 | 1/1/2:1 | | |
| 30995 | 1.0.26.2 (B) | 1/1/3:1 | | |
| 1.0.56.5/32 | | | SR | - |
| 3 | 1.0.56.5 | 1/1/2:1 | | |
| 10.20.1.1/32 | | | SR-OSPF-0 | - |
| 30010 | 1.0.26.2 | 1/1/3:1 | | |
| 40010 | 1.0.36.3 (B) | 1/1/4:1 | | |
| 10.20.1.2/32 | | | SR-OSPF-0 | - |
| 30011 | 1.0.26.2 | 1/1/3:1 | | |
| 50011/60001 | 1.0.56.5 (B) | 1/1/2:1 | | |
| 10.20.1.3/32 | | | SR-OSPF-0 | - |
| 40006 | 1.0.36.3 | 1/1/4:1 | | |
| 20006/30004 | 1.0.26.2 (B) | 1/1/3:1 | | |
| 10.20.1.4/32 | | | SR-OSPF-0 | - |
| 30007 | 1.0.26.2 | 1/1/3:1 | | |
| 60007 | 1.0.56.5 (B) | 1/1/2:1 | | |
| 10.20.1.5/32 | | | SR-OSPF-0 | - |
| 60008 | 1.0.56.5 | 1/1/2:1 | | |
| 50008/30001 | 1.0.26.2 (B) | 1/1/3:1 | | |

Total Entries : 14

=====

*A:Dut-F#

*A:Dut-C# show router fp-tunnel-table 1 10.20.1.5/32

=====

```

Tunnel Table Display
Legend:
B - FRR Backup
=====
Destination          NextHop          Intf/Tunnel      Protocol    Tunnel-ID
  Lbl
-----
10.20.1.5/32
  262135              10.10.5.5        2/1/1           LDP         -
  3                  10.20.1.5 (B)    SR
10.20.1.5/32
  474390              10.10.5.5        2/1/1           SR-ISIS-0   -
  474390/474389      10.10.12.2 (B)   lag-1
-----
Total Entries : 2
=====

```

icmp

- Syntax** `icmp [interface interface-name]`
- Context** `show>router`
- Description** This command displays Internet Control Message Protocol version 4 (ICMP) statistics. ICMP generates error messages (for example, ICMP destination unreachable messages) to report errors during processing and other diagnostic functions.
- Parameters** *interface-name* — specifies an existing IP interface name.
- Values** up to 32 characters
- Output** The following output is an example of router ICMP statistics, and [Table 16](#) describes the fields.

Sample Output

```

*A:cses-V93# show router icmp
=====
Global ICMP Stats
=====
Received
Total          : 0          Error          : 0
Destination Unreachable : 0          Redirect       : 0
Echo Request   : 0          Echo Reply     : 0
TTL Expired    : 0          Source Quench  : 0
Timestamp Request : 0        Timestamp Reply : 0
Address Mask Request : 0        Address Mask Reply : 0
Parameter Problem : 0
-----
Sent
Total          : 0          Error          : 0
Destination Unreachable : 0          Redirect       : 0
Echo Request   : 0          Echo Reply     : 0
TTL Expired    : 0          Source Quench  : 0

```

```

Timestamp Request      : 0          Timestamp Reply      : 0
Address Mask Request   : 0          Address Mask Reply   : 0
Parameter Problem      : 0
=====
*A:cses-V93# show router icmp interface "foo"
=====
Interface ICMP Stats
=====
Interface "foo"
-----
Received
Total                  : 0          Error                  : 0
Destination Unreachable : 0          Redirect              : 0
Echo Request           : 0          Echo Reply             : 0
TTL Expired            : 0          Source Quench          : 0
Timestamp Request      : 0          Timestamp Reply        : 0
Address Mask Request   : 0          Address Mask Reply     : 0
Parameter Problem      : 0
-----
Sent
Total                  : 0          Error                  : 0
Destination Unreachable : 0          Redirect              : 0
Echo Request           : 0          Echo Reply             : 0
TTL Expired            : 0          Source Quench          : 0
Timestamp Request      : 0          Timestamp Reply        : 0
Address Mask Request   : 0          Address Mask Reply     : 0
Parameter Problem      : 0
=====

```

Table 16 **ICMP Fields**

| Label | Description |
|-------------------------|--|
| Total | The total number of all messages. |
| Error | The number of error messages. |
| Destination Unreachable | The number of message that did not reach the destination. |
| Redirect | The number of packet redirects. |
| Echo Request | The number of echo requests. |
| Echo Reply | The number of echo replies. |
| TTL Expired | The number of messages that exceeded the time to live threshold. |
| Source Quench | The number of source quench requests (deprecated). |
| Timestamp Request | The number of timestamp requests. |
| Timestamp Reply | The number of timestamp replies. |
| Address Mask Request | The number of address mask requests (deprecated). |

Table 16 ICMP Fields (Continued)

| Label | Description |
|--------------------|--|
| Address Mask Reply | The number of address mask replies (deprecated). |
| Parameter Problem | The number of packets with a parameter problem in the IP header. |

icmp6

| | |
|--------------------|--|
| Syntax | icmp6 [interface <i>interface-name</i>] |
| Context | show>router |
| Description | This command displays Internet Control Message Protocol Version 6 (ICMPv6) statistics. ICMP generates error messages (for example, ICMP destination unreachable messages) to report errors during processing and other diagnostic functions. ICMPv6 packets can be used in the neighbor discovery protocol and path MTU discovery. |
| Parameters | <i>interface-name</i> — specifies an existing IP interface name. Values up to 32 characters |
| Output | The following output is an example of router ICMPv6 statistics, and Table 17 describes the fields. |

Sample Output

```
A:SR-3# show router icmp6
=====
Global ICMPv6 Stats
=====
Received
Total                : 0                Errors                : 0
Destination Unreachable : 0                Redirects                : 0
Time Exceeded         : 0                Pkt Too Big              : 0
Echo Request          : 0                Echo Reply                : 0
Router Solicits        : 0                Router Advertisements     : 0
Neighbor Solicits      : 0                Neighbor Advertisements   : 0
Parameter Problem      : 0
-----
Sent
Total                : 2                Errors                : 0
Destination Unreachable : 0                Redirects                : 0
Time Exceeded         : 0                Pkt Too Big              : 0
Echo Request          : 0                Echo Reply                : 0
Router Solicits        : 0                Router Advertisements     : 0
Neighbor Solicits      : 2                Neighbor Advertisements   : 0
Parameter Problem      : 0
=====
A:SR-3#
A:SR-3# show router icmp6 interface "foo"
=====
Interface ICMPv6 Stats
```

```

=====
Interface "foo"
-----
Received
Total          : 0          Errors          : 0
Destination Unreachable : 0          Redirects        : 0
Time Exceeded   : 0          Pkt Too Big      : 0
Echo Request    : 0          Echo Reply       : 0
Router Solicits : 0          Router Advertisements : 0
Neighbor Solicits : 0        Neighbor Advertisements : 0
Parameter Problem : 0
-----
Sent
Total          : 2          Errors          : 0
Destination Unreachable : 0          Redirects        : 0
Time Exceeded   : 0          Pkt Too Big      : 0
Echo Request    : 0          Echo Reply       : 0
Router Solicits : 0          Router Advertisements : 0
Neighbor Solicits : 2        Neighbor Advertisements : 0
Parameter Problem : 0
=====
A:SR-3#

```

Table 17 **ICMPv6 Fields**

| Label | Description |
|-------------------------|--|
| Total | The total number of all messages. |
| Destination Unreachable | The number of message that did not reach the destination. |
| Time Exceeded | The number of messages that exceeded the time threshold. |
| Echo Request | The number of echo requests. |
| Router Solicits | The number of times the local router was solicited. |
| Neighbor Solicits | The number of times the neighbor router was solicited. |
| Errors | The number of error messages. |
| Redirects | The number of packet redirects. |
| Pkt Too big | The number of packets that exceed appropriate size. |
| Echo Reply | The number of echo replies. |
| Router Advertisements | The number of times the router advertised its location. |
| Neighbor Advertisements | The number of times the neighbor router advertised its location. |
| Parameter Problem | The number of packets with a parameter problem in the IP header. |

if-attribute

| | |
|--------------------|--|
| Syntax | if-attribute |
| Context | show>router |
| Description | This command enables the context to display interface attribute related information. |

srlg-group

| | |
|--------------------|---|
| Syntax | srlg-group [<i>name</i>] |
| Context | show>router>if-attribute>srlg-group |
| Description | This command displays SRLG statistics. |
| Parameters | <i>name</i> — only displays entries associated with the specified SRLG name |
| Output | The following output is an example of SRLG statistics, and Table 18 describes the fields. |

Sample Output

```

B:CORE2# show router if-attribute srlg-group
=====
Interface Srlg Groups
=====
Group Name          Group Value  Penalty Weight
-----
1                    1           100
2                    2           200
3                    3           300
-----
No. of Groups: 3
=====
B:CORE2#

```

Table 18 **SRLG Fields**

| Label | Description |
|----------------|--|
| Group Name | The name of the SRLG. |
| Group Value | The integer value of the SRLG. |
| Penalty Weight | The penalty weight that is assigned to the SRLG. |
| No. of Groups | The total number of displayed SRLGs. |

interface

| | |
|--------------------|---|
| Syntax | interface [<i>interface-name</i>] |
| Context | show>router>icmpv6 |
| Description | This command displays interface ICMPv6 statistics. |
| Parameters | <i>interface-name</i> — only displays entries associated with the specified IP interface name |
| Output | The following output is an example of ICMPv6 interface statistics, and Table 19 describes the fields. |

Sample Output

```

B:CORE2# show router icmp6 interface net1_1_2
=====
Interface ICMPv6 Stats
=====
Interface "net1_1_2"
-----
Received
Total                : 41                Errors                : 0
Destination Unreachable : 0                Redirects                : 0
Time Exceeded         : 0                Pkt Too Big             : 0
Echo Request          : 0                Echo Reply               : 0
Router Solicits        : 0                Router Advertisements    : 0
Neighbor Solicits      : 20               Neighbor Advertisements  : 21
-----
Sent
Total                : 47                Errors                : 0
Destination Unreachable : 0                Redirects                : 0
Time Exceeded         : 0                Pkt Too Big             : 0
Echo Request          : 0                Echo Reply               : 0
Router Solicits        : 0                Router Advertisements    : 0
Neighbor Solicits      : 27               Neighbor Advertisements  : 20
=====
B:CORE2#

```

Table 19 ICMP6 Interface Fields

| Label | Description |
|-------------------------|---|
| Total | The total number of all messages. |
| Destination Unreachable | The number of message that did not reach the destination. |
| Time Exceeded | The number of messages that exceeded the time threshold. |
| Echo Request | The number of echo requests. |
| Router Solicits | The number of times the local router was solicited. |

Table 19 ICMP6 Interface Fields (Continued)

| Label | Description |
|-------------------------|--|
| Neighbor Solicits | The number of times the neighbor router was solicited. |
| Errors | The number of error messages. |
| Redirects | The number of packet redirects. |
| Pkt Too big | The number of packets that exceed appropriate size. |
| Echo Reply | The number of echo replies. |
| Router Advertisements | The number of times the router advertised its location. |
| Neighbor Advertisements | The number of times the neighbor router advertised its location. |

interface

Syntax **interface** {{*ip-address* | *ip-int-name*} [**detail**] [*family*] | **summary** | **exclude-services**}
interface {*ip-address* | *ip-int-name*} **eth-cfm** [**detail**]
interface {*ip-address* | *ip-int-name*} **mac** [*ieee-address*]
interface {*ip-address* | *ip-int-name*} **statistics**
interface {*ip-address* | *ip-int-name*} **dist-cpu-protection** [**detail**]
interface {*ip-address* | *ip-int-name*} **policy-accounting** [*class* [*index*]]

Context show>router

Description This command displays the router IP interface table sorted by interface index.

Parameters *ip-address* — displays the interface information associated with the specified IP address

Values

ipv4-address a.b.c.d (host bits must be 0)
ipv6-address x:x:x:x:x:x:x (eight 16-bit
 pieces)
 x:x:x:x:x:d.d.d.d

x: [0 to FFFF]H
d: [0 to 255]D

ip-int-name — displays the interface information associated with the specified IP interface name. The name can be up to 32 characters in length.

summary — displays summary IP interface information for the router

detail — displays detailed IP interface information

exclude-services — displays IP interface information, excluding IP interfaces configured for customer services. Only core network IP interfaces are displayed.

family — specifies the router IP interface family to display

Values **ipv4** — displays only those peers that have the IPv4 family enabled

Values **ipv6** — displays the peers that are IPv6-capable

eth-cfm — displays Ethernet CFM information

mac — displays information associated with the MAC address

ieee-address — displays the information associated with the specified IEEE address.
The address is in the xx-xx-xx-xx-xx-xx or xx:xx:xx:xx:xx:xx format.

dist-cpu-protection — displays the Distributed CPU Protection parameters and status at the interface level

class — indicates whether to display accounting policy statistics for the source or destination class

Values source-class, dest-class

index — specifies an integer value for the accounting source or destination class index

Values 1 to 255

statistics — displays packet statistics for an interface on the router



Note: The **show router interface statistics** command also shows the MPLS statistics that are shown in using the **show router mpls interface statistics** command. This allows the operator to see MPLS statistics from interfaces that are not added to MPLS, such as a carrier's network interfaces. [Sample Output](#) for an example of the MPLS fields that are displayed. These fields are displayed regardless of the state of MPLS.

Output **Standard IP Interface Output**—The following output is an example of standard IP interface information, and [Table 20](#) describes the fields.

Sample Output

```
*A:mlstp-dutA# show router interface "AtoB_1"
=====
Interface Table (Router: Base)
=====
Interface-Name      Adm      Opr (v4/v6)  Mode      Port/SapId
IP-Address          PfxState
-----
AtoB_1              Down     Down/--      Network  1/2/3:1
Unnumbered If[system]              n/a
-----
Interfaces : 1

A:ALA-A# show router interface
=====
```

```
Interface Table (Router: Base)
=====
Interface-Name      Adm(v4/v6)  Opr (v4/v6)  Mode    Port/SapId
IP-Address          PfxState
-----
ip-100.0.0.2        Up/Up       Up/Up        Network lag-1
100.0.0.2/10        n/a
3FFE:1::2/64        PREFERRED
FE80::200:FF:FE00:4/64 PREFERRED
ip-100.128.0.2      Up/Up       Up/Up        Network lag-2
100.128.0.2/10      n/a
3FFE:2::2/64        PREFERRED
FE80::200:FF:FE00:4/64 PREFERRED
ip-24.2.4.4         Up/Up       Up/Up        Network 6/2/14
24.2.4.4/24         n/a
3FFE::1802:404/120  PREFERRED
FE80::200:FF:FE00:4/64 PREFERRED
system              Up/Up       Up/Up        Network system
200.200.200.4/32    n/a
3FFE::C8C8:C804/128 PREFERRED
-----
Interfaces : 4
=====
A:ALA-A#

A:ALA-A# show router interface 10.10.0.3/32
=====
Interface Table
=====
Interface-Name      Type IP-Address      Adm   Opr   Mode
-----
system              Pri  10.10.0.3/32    Up    Up    Network
=====
A:ALA-A#

A:ALA-A# show router interface exclude-services
=====
Interface Table
=====
Interface-Name      Type IP-Address      Adm   Opr   Mode
-----
system              Pri  10.10.0.3/32    Up    Up    Network
to-ser1             Pri  10.10.13.3/24    Up    Up    Network
to-ser4             Pri  10.10.34.3/24    Up    Up    Network
to-ser5             Pri  10.10.35.3/24    Up    Up    Network
to-ser6             n/a  n/a             Up    Down  Network
management          Pri  192.168.2.93/20  Up    Up    Network
=====
A:ALA-A#
```

Table 20 Standard IP Interface Field Descriptions

| Label | Description |
|----------------|------------------------|
| Interface-Name | The IP interface name. |

Table 20 Standard IP Interface Field Descriptions (Continued)

| Label | Description |
|-------------|--|
| Type | n/a — No IP address has been assigned to the IP interface, so the IP address type is not applicable. Pri — The IP address for the IP interface is the Primary address on the IP interface. Sec — The IP address for the IP interface is a secondary address on the IP interface. |
| IP-Address | The IP address and subnet mask length of the IP interface. n/a — Indicates no IP address has been assigned to the IP interface. |
| Adm | Down — The IP interface is administratively disabled. Up — The IP interface is administratively enabled. |
| Opr | Down — The IP interface is operationally disabled. Up — The IP interface is operationally disabled. |
| Mode | Network — The IP interface is a network/core IP interface. Service — The IP interface is a service IP interface. |
| Port/SAP Id | The physical network port or the SAP identifier associated with the IP interface. |

Detailed IP Interface Output — The following output is an example of detailed IP interface information, and [Table 21](#) describes the fields.

Sample Output

```
*A:Dut-C# show service id 10 interface "foo" detail
=====
Interface Table
=====
Interface
-----
If Name           : foo
Admin State       : Up
Oper (v4/v6)      : Up/Up
Protocols         : None
IP Addr/mask      : 1.2.3.3/24
Address Type      : Primary
IGP Inhibit       : Disabled
Broadcast Address : Host-ones
HoldUp-Time       : 0
Track Srrp Inst   : 0
IPv6 Address      : 3ffe::102:303/120
IPv6 Addr State   : PREFERRED
CGA modifier      : (Not Specified)
HoldUp-Time       : 0
Track Srrp Inst   : 0
Link Lcl Address  : fe80::200:ff:fe00:3/64
Link Lcl State    : PREFERRED
Ignore Port State : None
```

Details

| | | | |
|-------------------|-----------------------|--------------------|------------|
| Description | : (Not Specified) | | |
| If Index | : 29 | Virt. If Index | : 29 |
| Last Oper Chg | : 06/07/2016 15:02:00 | Global If Index | : 365 |
| Mon Oper Grp | : None | | |
| Srrp En Rtnng | : Disabled | Hold time | : N/A |
| SAP Id | : 2/1/1:10 | | |
| TOS Marking | : Trusted | If Type | : VPRN |
| SNTP B.Cast | : False | | |
| MAC Address | : 00:00:00:00:00:03 | Mac Accounting | : Disabled |
| Ingress stats | : Disabled | IPv6 DAD | : Enabled |
| TCP MSS V4 | : 0 | TCP MSS V6 | : 0 |
| ARP Timeout | : 14400s | IPv6 Nbr ReachTime | : 30s |
| ARP Retry Timer | : 5000ms | IPv6 stale time | : 14400s |
| ARP Limit | : Disabled | IPv6 Nbr Limit | : Disabled |
| ARP Threshold | : Disabled | IPv6 Nbr Threshold | : Disabled |
| ARP Limit Log On* | : Disabled | IPv6 Nbr Log Only | : Disabled |
| IP MTU | : (default) | | |
| IP Oper MTU | : 1500 | | |
| ARP Populate | : Disabled | Host Conn Verify | : Disabled |
| SHCV pol IPv4 | : None | | |
| Cflowd (unicast) | : None | Cflowd (multicast) | : None |
| LdpSyncTimer | : None | | |
| LSR Load Balance | : system | | |
| EGR Load Balance | : both | | |
| Vas If Type | : none | | |
| TEID Load Balance | : Disabled | | |
| SPI Load Balance | : Disabled | | |
| uRPF Chk | : disabled | | |
| uRPF Ipv6 Chk | : disabled | | |
| PTP HW Assist | : Disabled | | |
| Rx Pkts | : 87 | Rx Bytes | : 6216 |
| Rx V4 Pkts | : N/A | Rx V4 Bytes | : N/A |
| Rx V6 Pkts | : N/A | Rx V6 Bytes | : N/A |
| Tx Pkts | : 42 | Tx Bytes | : 3612 |
| Tx V4 Pkts | : 0 | Tx V4 Bytes | : 0 |
| Tx V4 Discard Pk* | : 0 | Tx V4 Discard Byt* | : 0 |
| Tx V6 Pkts | : 42 | Tx V6 Bytes | : 3612 |
| Tx V6 Discard Pk* | : 0 | Tx V6 Discard Byt* | : 0 |
| Mpls Rx Pkts | : 0 | Mpls Rx Bytes | : 0 |
| Mpls Tx Pkts | : 0 | Mpls Tx Bytes | : 0 |
| Proxy ARP Details | | | |
| Rem Proxy ARP | : Disabled | Local Proxy ARP | : Disabled |
| Policies | : none | | |

Proxy Neighbor Discovery Details

Local Pxy ND : Disabled
Policies : none

Secure ND Details

Secure ND : Disabled

DHCP no local server

DHCP Details

Description : (Not Specified)

| | | | |
|--|------------|-------------------|------------|
| Admin State | : Down | Lease Populate | : 0 |
| Gi-Addr | : 1.2.3.3* | Gi-Addr as Src Ip | : Disabled |
| * = inferred gi-address from interface | | IP address | |
| Action | : Keep | Trusted | : Disabled |

```

DHCP Proxy Details
Admin State      : Down
Lease Time       : N/A
Emul. Server     : Not configured
Subscriber Authentication Details
Auth Policy      : None
DHCP6 Relay Details
Description      : (Not Specified)
Admin State      : Down
Oper State       : Down
If-Id Option     : None
Src Addr         : Not configured
Python plcy      : (Not Specified)
Lease Populate   : 0
Nbr Resolution   : Disabled
Remote Id        : Disabled
DHCP6 Server Details
Admin State      : Down
Max. Lease States : 8000
ISA Tunnel redundant next-hop information
Static Next-Hop :
Dynamic Next-Hop :
ICMP Details
Redirects        : Number - 100      Time (seconds) - 10
Unreachables     : Number - 100      Time (seconds) - 10
TTL Expired      : Number - 100      Time (seconds) - 10
Parameter Problem: Number - 100      Time (seconds) - 10
ICMP Mask Reply  : True
ICMPv6 Details
Packet Too Big   : Number - 100      Time (seconds) - 10
Parameter Problem: Number - 100      Time (seconds) - 10
Redirects        : Number - 100      Time (seconds) - 10
Time Exceeded    : Number - 100      Time (seconds) - 10
Unreachables     : Number - 100      Time (seconds) - 10
IPCP Address Extension Details
Peer IP Addr     : Not configured
Peer Pri DNS Addr: Not configured
Peer Sec DNS Addr: Not configured
-----
Admin Groups
-----
No Matching Entries
-----
Srlg Groups
-----
No Matching Entries
-----
QoS Queue-Group Redirection Details
-----
Ingress FP QGrp  : (none)      Egress Port QGrp  : (none)
Ing FP QGrp Inst : (none)      Egr Port QGrp Inst: (none)
-----
Interfaces : 1
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-C#

*A:Dut-C>config>router>mpls# /show router 1 interface "To_B_1" detail
=====

```



```

Interface Table (Service: 1)
=====
-----
Interface
-----
If Name       : To_B_1
Admin State   : Up                               Oper (v4/v6)   : Down/Down
Down Reason Code : assocObjNotReady
Down Reason V4  : assocObjNotReady
Down Reason V6  : assocObjNotReady ifProtoOperDown
Protocols      : OSPFv2
IP Addr/mask    : 11.11.11.1/24                 Address Type   : Primary
IGP Inhibit     : Disabled                       Broadcast Address : Host-ones
HoldUp-Time     : 0                             Track Srrp Inst : 0
-----
Details
-----
Description    : (Not Specified)
If Index       : 5                               Virt. If Index : 5
Last Oper Chg  : 07/21/2016 21:46:23           Global If Index : 258
Mon Oper Grp   : None
Srrp En Rtng   : Disabled                       Hold time      : N/A
SDP Id         : spoke-230:1
Spoke-SDP Details
Admin State    : Up                               Oper State     : Down
Hash Label     : Disabled                       Hash Lbl Sig Cap : Disabled
Oper Hash Label : Disabled
Entropy Label   : Disabled
Peer Fault Ip   : None
Local Pw Bits   : pwNotForwarding
Peer Pw Bits    : None
Peer Vccv CV Bits : lspPing bfdFaultDet
Peer Vccv CC Bits : mplsRouterAlertLabel
Flags          : LabelStackLimitExceeded
TOS Marking     : Trusted                       If Type        : VPRN
SNTP B.Cast     : False
MAC Address     : 0e:86:ff:00:00:00             Mac Accounting  : Disabled
Ingress stats   : Disabled                       IPv6 DAD        : Enabled
TCP MSS V4      : 0                             TCP MSS V6      : 0
ARP Timeout     : 14400s                         IPv6 Nbr ReachTime: 30s
ARP Retry Timer : 5000ms                         IPv6 stale time  : 14400s
ARP Limit       : Disabled                       IPv6 Nbr Limit   : Disabled
ARP Threshold   : Disabled                       IPv6 Nbr Threshold: Disabled
ARP Limit Log On* : Disabled                     IPv6 Nbr Log Only : Disabled
IP MTU          : 1500
IP Oper MTU     : 0
ARP Populate    : Disabled                       Host Conn Verify : Disabled
SHCV pol IPv4   : None
Cflowd (unicast) : None                         Cflowd (multicast): None
LdpSyncTimer    : None
LSR Load Balance : system
EGR Load Balance : both
Vas If Type     : none
TEID Load Balance: Disabled
SPI Load Balance : Disabled
uRPF Chk        : disabled
uRPF Ipv6 Chk   : disabled
PTP HW Assist   : Disabled
Rx Pkts         : 0                             Rx Bytes        : 0

```

```

Rx V4 Pkts      : N/A
Rx V6 Pkts      : N/A
Tx Pkts         : 22
Tx V4 Pkts      : 0
Tx V4 Discard Pk*: 0
Tx V6 Pkts      : 0
Tx V6 Discard Pk*: 0
Mpls Rx Pkts    : 0
Mpls Tx Pkts    : 0
Proxy ARP Details
Rem Proxy ARP    : Disabled
Policies         : none
Proxy Neighbor Discovery Details
Local Pxy ND     : Disabled
Policies         : none
DHCP no local server
DHCP Details
Description      : (Not Specified)
Admin State      : Down
Gi-Addr          : 11.11.11.1*
* = inferred gi-address from interface IP address
Action           : Keep
Lease Populate   : 0
Gi-Addr as Src Ip : Disabled
Trusted          : Disabled
DHCP Proxy Details
Admin State      : Down
Lease Time       : N/A
Emul. Server     : Not configured
Subscriber Authentication Details
Auth Policy      : None
DHCP6 Relay Details
Description      : (Not Specified)
Admin State      : Down
Oper State       : Down
If-Id Option     : None
Src Addr         : Not configured
Python plcy      : (Not Specified)
Lease Populate   : 0
Nbr Resolution   : Disabled
Remote Id        : Disabled
DHCP6 Server Details
Admin State      : Down
Max. Lease States : 8000
ISA Tunnel redundant next-hop information
Static Next-Hop  :
Dynamic Next-Hop :
ICMP Details
Redirects        : Number - 100
Unreachables     : Number - 100
TTL Expired      : Number - 100
Parameter Problem: Number - 100
ICMP Mask Reply  : True
Time (seconds)   - 10
ICMPv6 Details
Packet Too Big   : Number - 100
Parameter Problem: Number - 100
Redirects        : Number - 100
Time Exceeded    : Number - 100
Unreachables     : Number - 100
Time (seconds)   - 10
IPCP Address Extension Details
Peer IP Addr     : Not configured
Peer Pri DNS Addr: Not configured
Peer Sec DNS Addr: Not configured
-----
Admin Groups
-----

```

```

No Matching Entries
-----
Srlg Groups
-----
No Matching Entries
-----
QoS Queue-Group Redirection Details
-----
Ingress FP QGrp  : (none)          Egress Port QGrp  : (none)
Ing FP QGrp Inst : (none)          Egr Port QGrp Inst: (none)
=====
* indicates that the corresponding row element may have been truncated.

```

Table 21 Detailed IP Interface Field Descriptions

| Label | Description |
|------------------|---|
| If Name | The IP interface name. |
| Admin State | Down — The IP interface is administratively disabled. Up — The IP interface is administratively enabled. |
| Oper State | Down — The IP interface is operationally disabled. Up — The IP interface is operationally enabled. |
| IP Addr/mask | The IP address and subnet mask length of the IP interface. Not Assigned — Indicates no IP address has been assigned to the IP interface. |
| IPv6 Addr | The IPv6 address of the interface. |
| If Index | The interface index of the IP router interface. |
| Virt If Index | The virtual interface index of the IP router interface. |
| Last Oper Change | The last change in operational status. |
| Global If Index | The global interface index of the IP router interface. |
| Sap ID | The SAP identifier. |
| TOS Marker | The TOS byte value in the logged packet. |
| If Type | Network — The IP interface is a network/core IP interface. Service — The IP interface is a service IP interface. |
| SNTP B.cast | Displays if the broadcast-client global parameter is configured. |
| IES ID | The IES identifier. |
| QoS Policy | The QoS policy ID associated with the IP interface. |
| MAC Address | The MAC address of the interface. |

Table 21 Detailed IP Interface Field Descriptions (Continued)

| Label | Description |
|------------------|---|
| Arp Timeout | The ARP timeout for the interface, in seconds, which is the time an ARP entry is maintained in the ARP cache without being refreshed. |
| ICMP Mask Reply | False — The IP interface will not reply to a received ICMP mask request. True — The IP interface will reply to a received ICMP mask request. |
| Arp Populate | Displays whether ARP is enabled or disabled. |
| Host Conn Verify | The host connectivity verification. |
| LdpSyncTimer | Specifies the IGP/LDP sync timer value. |
| uRPF Chk | Specifies whether unicast RPF (uRPF) Check is enabled on this interface. |
| uRPF Iv6 Chk | Specifies whether unicast RPF (uRPF) Check IPv6 is enabled on this interface. |
| PTP HW Assist | Specifies whether the PTP Hardware Assist function is enabled on this interface. |
| cflowd | Specifies the type of cflowd analysis that is applied to the interface. acl — ACL cflowd analysis is applied to the interface. interface — Interface cflowd analysis is applied to the interface. none — No cflowd analysis is applied to the interface. |

Statistics IP Interface Output — The following output is an example of router IP interface statistics when **enable-interface-statistics** is enabled, and [Table 22](#) describes the fields.

Sample Output

```
A:ALA-A# show router interface "to_ixia" statistics
=====
Interface Statistics
=====
If Name           : to_Ixia
Admin State       : Up
Oper (v4/v6)      : Up/Up
Rx Pkts           : 6244
Rx Bytes          : 599424
Rx V4 Pkts        : 3122
Rx V4 Bytes       : 299712
Rx V6 Pkts        : 3122
Rx V6 Bytes       : 299712
Tx Pkts           : 0
Tx Bytes          : 0
Tx V4 Pkts        : 0
Tx V4 Bytes       : 0
Tx V4 Discard Pk* : 0
Tx V4 Discard Byt* : 0
Tx V6 Pkts        : 0
Tx V6 Bytes       : 0
```

```

Tx V6 Discard Pk*: 0
uRPF Chk Fail Pk*: 6244
uRPF Fail V4 Pk : 3122
uRPF Fail V6 Pk : 3122
Mpls Rx Pkts : 0
Mpls Tx Pkts : 0

Tx V6 Discard Byt*: 0
uRPF Fail Bytes : 487032
uRPF Fail V4 Byt : 243516
uRPF Fail V6 Byt : 243516
Mpls Rx Bytes : 0
Mpls Tx Bytes : 0
=====

```

Table 22 Statistics IP Interface Fields

| Label | Description |
|-------------|--|
| Ifname | The interface name. |
| Admin State | The administrative status of the router interface. |
| Oper | The operational status of the router instance. |

Summary IP Interface Output — The following output is an example of summary IP information, and [Table 23](#) describes the fields.

Sample Output

```

A:ALA-A# show router interface summary
=====
Router Summary (Interfaces)
=====
Instance  Router Name                Interfaces  Admin-Up  Oper-Up
-----
1         Base                      7          7          5
=====

```

Table 23 Summary IP Interface Fields

| Label | Description |
|-------------|--|
| Instance | The router instance number. |
| Router Name | The name of the router instance. |
| Interfaces | The number of IP interfaces in the router instance. |
| Admin-Up | The number of administratively enabled IP interfaces in the router instance. |
| Oper-Up | The number of operationally enabled IP interfaces in the router instance. |

routes

| | |
|--------------------|--|
| Syntax | routes alternative |
| Context | show:router>isis |
| Description | This command displays IS-IS route information. |
| Output | The following output is an example of IS-IS route information. |

Sample Output

```
*A:SRR# show router isis routes 1.1.1.0/24
=====
Route Table
=====
Prefix[Flags]                Metric    Lvl/Typ   Ver.    SysID/Hostname
  NextHop                    MT        AdminTag
-----
1.1.1.0/24 [L]                7540      1/Int.    6109    SRL
  60.60.1.1                    0         0
-----

No. of Routes: 1
Flags: L = LFA nexthop available
=====

*A:SRR#
*A:SRR# show router isis routes 1.1.1.0/24 alternative
=====
Route Table
=====
Prefix[Flags]                Metric    Lvl/Typ   Ver.    SysID/Hostname
  NextHop                    MT        AdminTag
Alt-Nexthop                  Alt-Metric Alt-Type
-----
1.1.1.0/24                    7550      1/Int.    6114    SRL
  60.60.1.1                    0         0
  11.22.12.4 (LFA)             16784764  linkProtection
-----

No. of Routes: 1
Flags: LFA = Loop-Free Alternate nexthop
=====

*A:SRR#

*A:Dut-B# show router isis routes
=====
Route Table
=====
Prefix [Flags]                Metric    Lvl/Typ   Ver.    SysID/Hostname
  NextHop                    MT        AdminTag
-----
10.20.1.2/32                  0         1/Int.    3       Dut-B
  0.0.0.0                      0         0
10.20.1.3/32 [L]              10        2/Int.    2       Dut-C
  10.20.3.3                    0         0
10.20.1.4/32                  10        2/Int.    3       Dut-D
  10.20.4.4                    0         0
10.20.1.5/32                  20        2/Int.    3       Dut-C
```

```

10.20.3.3          0          0
10.20.1.6/32       20        2/Int.   3      Dut-D
10.20.4.4          0          0
10.20.3.0/24       10        1/Int.   3      Dut-B
0.0.0.0            0          0
10.20.4.0/24       10        1/Int.   3      Dut-B
0.0.0.0            0          0
10.20.5.0/24       20        2/Int.   2      Dut-C
10.20.3.3          0          0
10.20.6.0/24       20        2/Int.   4      Dut-D
10.20.4.4          0          0
10.20.9.0/24       20        2/Int.   3      Dut-D
10.20.4.4          0          0
10.20.10.0/24      30        2/Int.   3      Dut-C
10.20.3.3          0          0

```

Routes : 11

Flags: L = LFA nexthop available

=====

*A:Dut-B#

*A:Dut-B# show router isis routes alternative

Route Table

```

=====
Prefix [Flags]          Metric   Lvl/Typ   Ver.   SysID/Hostname
NextHop                MT        AdminTag
Alt-Nexthop            Alt-Metric
-----
10.20.1.2/32            0        1/Int.    3      Dut-B
0.0.0.0                 0        0
10.20.1.3/32            10       2/Int.    2      Dut-C
10.20.3.3               0        0
10.20.3.3 (lfa)        15
10.20.1.4/32            10       2/Int.    3      Dut-D
10.20.4.4               0        0
10.20.1.5/32            20       2/Int.    3      Dut-C
10.20.3.3               0        0
10.20.1.6/32            20       2/Int.    3      Dut-D
10.20.4.4               0        0
10.20.3.0/24            10       1/Int.    3      Dut-B
0.0.0.0                 0        0
10.20.4.0/24            10       1/Int.    3      Dut-B
0.0.0.0                 0        0
10.20.5.0/24            20       2/Int.    2      Dut-C
10.20.3.3               0        0
10.20.6.0/24            20       2/Int.    4
4      Dut-D
10.20.4.4               0        0
10.20.9.0/24            20       2/Int.    3      Dut-D
10.20.4.4               0        0
10.20.10.0/24           30       2/Int.    3      Dut-C
10.20.3.3               0        0

```

Routes : 11

Flags: LFA = Loop-Free Alternate nexthop

=====

*A:Dut-B#

bindings

| | |
|--------------------|---|
| Syntax | bindings active |
| Context | show>router>ldp |
| Description | This command displays LDP bindings information. |
| Output | The following output is an example of LDP bindings information. |

Sample Output

```
*A:Dut-A# show router ldp bindings active
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       WP - Label Withdraw Pending, BU - Alternate For Fast Re-Route
       (S) - Static      (M) - Multi-homed Secondary Support
       (B) - BGP Next Hop (BU) - Alternate Next-hop for Fast Re-Route
=====
LDP IPv4 Prefix Bindings (Active)
=====
Prefix                Op    IngLbl    EgrLbl    EgrIntf/LspId  EgrNextHop
-----
10.20.1.1/32          Pop   131071    --        --             --
10.20.1.2/32          Push  --        131071    1/1/1          10.10.1.2
10.20.1.2/32          Swap 131070    131071    1/1/1          10.10.1.2
10.20.1.2/32          Push  --        262141BU  1/1/2          10.10.2.3
10.20.1.2/32          Swap 131070    262141BU  1/1/2          10.10.2.3
10.20.1.3/32          Push  --        131069BU  1/1/1          10.10.1.2
10.20.1.3/32          Swap 131069    131069BU  1/1/1          10.10.1.2
10.20.1.3/32          Push  --        262143    1/1/2          10.10.2.3
10.20.1.3/32          Swap 131069    262143    1/1/2          10.10.2.3
10.20.1.4/32          Push  --        131068    1/1/1          10.10.1.2
10.20.1.4/32          Swap 131068    131068    1/1/1          10.10.1.2
10.20.1.4/32          Push  --        262140BU  1/1/2          10.10.2.3
10.20.1.4/32          Swap 131068    262140BU  1/1/2          10.10.2.3
10.20.1.5/32          Push  --        131067BU  1/1/1          10.10.1.2
10.20.1.5/32          Swap 131067    131067BU  1/1/1          10.10.1.2
10.20.1.5/32          Push  --        262139    1/1/2          10.10.2.3
10.20.1.5/32          Swap 131067    262139    1/1/2          10.10.2.3
10.20.1.6/32          Push  --        131066    1/1/1          10.10.1.2
10.20.1.6/32          Swap 131066    131066    1/1/1          10.10.1.2
10.20.1.6/32          Push  --        262138BU  1/1/2          10.10.2.3
10.20.1.6/32          Swap 131066    262138BU  1/1/2          10.10.2.3
-----
No. of IPv4 Prefix Active Bindings: 10
=====
LDP IPv6 Prefix Bindings (Active)
=====
Prefix                Op    IngLbl    EgrLbl
EgrNextHop            EgrIf/LspId
-----
No Matching Entries Found
=====
```



```

LDP Generic IPv4 P2MP Bindings (Active)
=====
P2MP-Id                               Interface
RootAddr                             Op           IngLbl    EgrLbl
EgrNH                                EgrIf/LspId
-----
No Matching Entries Found
=====

LDP Generic IPv6 P2MP Bindings (Active)
=====
P2MP-Id                               Interface
RootAddr                             Op           IngLbl    EgrLbl
EgrNH                                EgrIf/LspId
-----
No Matching Entries Found
=====

LDP In-Band-SSM IPv4 P2MP Bindings (Active)
=====
Source
Group                                Interface
RootAddr                             Op           IngLbl    EgrLbl
EgrNH                                EgrIf/LspId
-----
No Matching Entries Found
=====

LDP In-Band-SSM IPv6 P2MP Bindings (Active)
=====
Source
Group                                Interface
RootAddr                             Op           IngLbl    EgrLbl
EgrNH                                EgrIf/LspId
-----
No Matching Entries Found
=====

LDP In-Band-VPN-SSM IPv4 P2MP Bindings (Active)
=====
Source
Group                                RD           Op
RootAddr                             Interface    IngLbl    EgrLbl
EgrNH                                EgrIf/LspId
-----
No Matching Entries Found
=====

LDP In-Band-VPN-SSM IPv6 P2MP Bindings (Active)
=====
Source
Group                                RD           Op
RootAddr                             Interface    IngLbl    EgrLbl
EgrNH                                EgrIf/LspId
-----
No Matching Entries Found
=====

*A:Dut-A# show router ldp bindings

```

```

=====
LDP Bindings (IPv4 LSR ID 1.1.1.1:0)
              (IPv6 LSR ID ::[0])
=====
Legend: U - Label In Use, N - Label Not In Use, W - Label Withdrawn
       S - Status Signaled Up, D - Status Signaled Down
       E - Epipe Service, V - VPLS Service, M - Mirror Service
       A - Apipe Service, F - Fpipe Service, I - IES Service, R - VPRN service
       P - Ipipe Service, WP - Label Withdraw Pending, C - Cpipe Service
       BU - Alternate For Fast Re-Route, TLV - (Type, Length: Value)
=====
LDP IPv4 Prefix Bindings
=====
Prefix          Peer          IngLbl      EgrLbl EgrIntf/  EgrNextHop
                  LspId
-----
10.20.1.1/32    10.20.1.2    131071U    --    --        --
10.20.1.1/32    10.20.1.3    131071U    --    --        --
10.20.1.2/32    10.20.1.2    --        131071 1/1/1    10.10.1.2
10.20.1.2/32    10.20.1.3    131070U    262141 1/1/2    10.10.2.3
10.20.1.3/32    10.20.1.2    131069U    131069 1/1/1    10.10.1.2
10.20.1.3/32    10.20.1.3    --        262143 1/1/2    10.10.2.3
10.20.1.4/32    10.20.1.2    131068N    131068 1/1/1    10.10.1.2
10.20.1.4/32    10.20.1.3    131068BU   262140 1/1/2    10.10.2.3
10.20.1.5/32    10.20.1.2    131067U    131067 1/1/1    10.10.1.2
10.20.1.5/32    10.20.1.3    131067N    262139 1/1/2    10.10.2.3
10.20.1.6/32    10.20.1.2    131066N    131066 1/1/1    10.10.1.2
10.20.1.6/32    10.20.1.3    131066BU   262138 1/1/2    10.10.2.3
-----
No. of IPv4 Prefix Bindings: 12
=====
LDP IPv6 Prefix Bindings
=====
Prefix          IngLbl      EgrLbl
Peer            EgrIntf/LspId
EgrNextHop
-----
No Matching Entries Found
=====
LDP Generic IPv4 P2MP Bindings
=====
P2MP-Id
RootAddr          Interface    IngLbl      EgrLbl
EgrNH             EgrIf/LspId
Peer
-----
100
1.1.1.1           Unknw      --          131051
90.90.90.2        1/1/6
2.2.2.2:0

104
1.1.1.1           Unknw      --          131050
90.90.90.2        1/1/6
2.2.2.2:0

600

```

| | | | |
|------------|-------|----|--------|
| 1.1.1.1 | Unknw | -- | 131049 |
| 90.90.90.2 | 1/1/6 | | |
| 2.2.2.2:0 | | | |
| 700 | | | |
| 1.1.1.1 | Unknw | -- | 131048 |
| 90.90.90.2 | 1/1/6 | | |
| 2.2.2.2:0 | | | |
| 800 | | | |
| 1.1.1.1 | Unknw | -- | 131047 |
| 90.90.90.2 | 1/1/6 | | |
| 2.2.2.2:0 | | | |
| 900 | | | |
| 1.1.1.1 | Unknw | -- | 131046 |
| 90.90.90.2 | 1/1/6 | | |
| 2.2.2.2:0 | | | |
| 1500 | | | |
| 1.1.1.1 | Unknw | -- | 131045 |
| 90.90.90.2 | 1/1/6 | | |
| 2.2.2.2:0 | | | |
| 100 | | | |
| 6.6.6.6 | Unknw | -- | 131044 |
| 90.90.90.2 | 1/1/6 | | |
| 2.2.2.2:0 | | | |
| 900 | | | |
| 6.6.6.6 | Unknw | -- | 131043 |
| 90.90.90.2 | 1/1/6 | | |
| 2.2.2.2:0 | | | |

No. of Generic IPv4 P2MP Bindings: 9
=====

LDP Generic IPv6 P2MP Bindings
=====

| P2MP-Id | Interface | IngLbl | EgrLbl |
|----------|-------------|--------|--------|
| RootAddr | EgrIf/LspId | | |
| EgrNH | | | |
| Peer | | | |

No Matching Entries Found
=====

LDP In-Band-SSM IPv4 P2MP Bindings
=====

| Source | Interface | IngLbl | EgrLbl |
|----------|-------------|--------|--------|
| Group | EgrIf/LspId | | |
| RootAddr | | | |
| EgrNH | | | |
| Peer | | | |

No Matching Entries Found
=====

LDP In-Band-SSM IPv6 P2MP Bindings

```

=====
Source
Group
RootAddr          Interface      IngLbl  EgrLbl
EgrNH             EgrIf/LspId
Peer
-----

```

No Matching Entries Found

LDP In-Band-VPN-SSM IPv4 P2MP Bindings

```

=====
Source
Group              RD
RootAddr          Interface      IngLbl  EgrLbl
EgrNH             EgrIf/LspId
Peer
-----

```

```

1.1.1.1
225.0.0.1          1.1.1.1:100
3.3.3.3            Unknwn           --        100
60.60.60.1         1/1/1
2.2.2.2:100

```

```

1.1.1.1
225.0.0.1          1.1.1.1:100
3.3.3.3            Unknwn           --        100
60.60.60.1         1/1/1
2.2.2.2:100

```

```

1.1.1.1
225.0.0.1          1.1.1.1:100
3.3.3.3            Unknwn           --        100
60.60.60.1         1/1/1
2.2.2.2:100

```

No. of In-Band-VPN-SSM IPv4 P2MP Bindings: 3

LDP In-Band-VPN-SSM IPv6 P2MP Bindings

```

=====
Source
Group              RD
RootAddr          Interface      IngLbl  EgrLbl
EgrNH             EgrIf/LspId
Peer
-----

```

```

1.1.1.1
225.0.0.1          1.1.1.1:100
2000::3000         Unknwn           --        100
60.60.60.1         1/1/1
2.2.2.2:100

```

```

1.1.1.1
225.0.0.1          1.1.1.1:100
2000::3000         Unknwn           --        100
60.60.60.1         1/1/1

```

```

2.2.2.2:100

1.1.1.1
225.0.0.1          1.1.1.1:100
2000::3000         Unknwn          --          100
60.60.60.1         1/1/1
2.2.2.2:100

-----
No. of In-Band-VPN-SSM IPv6 P2MP Bindings: 3
=====
LDP Service FEC 128 Bindings
=====
Type                VCId      SDPId      IngLbl  LMTU
Peer                SvcId                      EgrLbl  RMTU
-----
?-Eth                100      R. Src     --      None
2.2.2.2:0            Ukwn                      131023D 986

?-Eth                500      R. Src     --      None
2.2.2.2:0            Ukwn                      131022D 1386

?-Eth                2001     R. Src     --      None
2.2.2.2:0            Ukwn                      131019D 986

?-Eth                2003     R. Src     --      None
2.2.2.2:0            Ukwn                      131017D 986

?-Ipipe             1800     R. Src     --      None
2.2.2.2:0            Ukwn                      131014D 1486

-----
No. of VC Labels: 5
=====
LDP Service FEC 129 Bindings
=====
SAII                AGII      IngLbl  LMTU
TAII                Type      EgrLbl  RMTU
Peer                SvcId      SDPId
-----
No Matching Entries Found
=====

```

mvpn

| | |
|--------------------|---|
| Syntax | mvpn |
| Context | show>router router-instance |
| Description | This command displays multicast VPN related information. The router instance must be specified. |
| Output | The following output is an example of MVPN information for the router-instance. |

Sample Output

```
*A:Dut-C# show router 1 mvpn
=====
MVPN 1 configuration data
=====
signaling          : Bgp                auto-discovery      : Enabled
UMH Selection      : Highest-Ip          intersite-shared     : Enabled
vrf-import         : N/A
vrf-export         : N/A
vrf-target         : target:1:1
C-Mcast Import RT  : target:10.20.1.3:2

ipmsi              : pim-asm 224.1.1.1
admin status       : Up                  three-way-hello      : N/A
hello-interval     : N/A                  hello-multiplier     : 35 * 0.1
tracking support    : Disabled             Improved Assert      : N/A

spmsi              : pim-ssm 225.0.0.0/32
join-tlv-packing   : N/A
data-delay-interval: 3 seconds
data-threshold     : 224.0.0.0/4 --> 1 kbps
=====
```

neighbor

- Syntax** `neighbor [ip-int-name | ip-address | mac ieee-mac-address | summary]`
- Context** `show>router`
- Description** This command displays information about the IPv6 neighbor cache.
- Parameters**
- ip-int-name* — specifies the IP interface name.
 - ip-address* — specifies the address of the IPv6 interface address.
 - mac ieee-mac-address* — specifies the MAC address.
 - summary** — displays summary neighbor information.
- Output** **Neighbor Output** — The following output is an example of IPv6 neighbor information, and [Table 24](#) describes the fields.

Sample Output

```
B:CORE2# show router neighbor
=====
Neighbor Table (Router: Base)
=====
IPv6 Address          State      Interface      Type      RTR
MAC Address           Expiry
-----
FE80::203:FAFF:FE78:5C88    STALE      net1_1_2
00:16:4d:50:17:a3          03h52m08s  Dynamic      Yes
FE80::203:FAFF:FE81:6888    net1_2_3
```

```

00:03:fa:1a:79:22          STALE          03h29m28s      Dynamic      Yes
-----
No. of Neighbor Entries: 2
=====
B:CORE2#

```

Table 24 Neighbor Fields

| Label | Description |
|--------------|---|
| IPv6 Address | Displays the IPv6 address. |
| Interface | Displays the name of the IPv6 interface name. |
| MAC Address | Specifies the link-layer address. |
| State | Displays the current administrative state. |
| Exp | Displays the number of seconds until the entry expires. |
| Type | Displays the type of IPv6 interface. |
| Interface | Displays the interface name. |
| Rtr | Specifies whether a neighbor is a router. |
| Mtu | Displays the MTU size. |

network-domains

- Syntax** `network-domains [detail] [network-domain-name]`
- Context** `show>router`
- Description** This command displays network-domains information.
- Parameters** **detail** — displays detailed network-domains information
network-domain-name — displays information for a specific network domain
- Output** The following output is an example of network domain information.

Sample

```

*A:Dut-T>config>router# show router network-domains
=====
Network Domain Table
=====
Network Domain          Description
-----
net1                    Network domain 1
default                Default Network Domain
=====

```

```

Network Domains : 2
=====
*A:Dut-T>config>router#

*A:Dut-T>config>router# show router network-domains detail
=====
Network Domain Table (Router: Base)
=====
-----
Network Domain           : net1
-----
Description               : Network domain 1
No. Of Ifs Associated     : 2
No. Of SDPs Associated    : 0

-----
Network Domain           : default
-----
Description               : Default Network Domain
No. Of Ifs Associated     : 3
No. Of SDPs Associated    : 0
=====
*A:Dut-T>config>router#

*A:Dut-T>config>router# show router network-domains "net1" interface-association
=====
Interface Network Domain Association Table
=====
Interface Name           Port           Network Domain
-----
intf1                    1/2/2        net1
intf2                    6/1/2        net1
-----
Interfaces : 2
=====
*A:Dut-T>config>router#

*A:Dut-T>config>service# show router network-domains "net1" sdp-association
=====
SDP Network Domain Association Table
=====
SDP Id                   Network Domain
-----
100                      net1
-----
SDPs : 1
=====
*A:Dut-T>config>service#

```

origin-validation

Syntax **origin-validation**

| | |
|--------------------|--|
| Context | show>router |
| Description | This command enables the context to display origin validation information. |

database

| | |
|--------------------|--|
| Syntax | database [<i>family</i>] [<i>ip-prefix/ip-prefix-length</i>] [<i>upto prefix-length2</i>][<i>origin-as as-number</i>] database [<i>family</i>] [<i>ip-prefix/ip-prefix-length</i>] { <i>longer</i> } database { <i>summary</i> } database [<i>family</i>] [{ <i>static</i> }] |
| Context | show>router>origin-validation |
| Description | This command displays database information. |
| Parameters | <i>family</i> — specifies the type of routing information to be displayed <div style="margin-left: 40px;"> Values ipv4 — displays IPv4 entries ipv6 — displays IPv6 entries </div> <i>ip-prefix/ip-prefix-length</i> — displays routes only matching the specified IP address and length <div style="margin-left: 40px;"> Values 64 characters maximum length </div> <i>prefix-length2</i> — displays routes matching up to the specified length <div style="margin-left: 40px;"> Values 1 to 128 </div> origin-as as-number — specifies the origin AS number <div style="margin-left: 40px;"> Values <i>as-number</i> — 0 to 4294967295 </div> longer — displays routes matching the <i>ip-prefix-ip-prefix-length</i> and routes with longer masks summary — displays database summary information static — displays static routes |
| Output | The following output is an example of database information. |

Sample Output

```

A:Dut-C# show router origin-validation database
=====
Static and Dynamic VRP Database Entries
=====
Prefix Range [Flags]                               Origin AS
  Session IP [Flags]
-----
10.0.0.0/16-24 [Static-V]                           65001
-
172.16.0.0/12-12 [Dynamic]                          65002
  192.168.1.1 [B]

```

```
-----
No. of Vrp Database Entries: 2
-----
Flags: B = Base instance session
      M = Management instance session
      Static-V = Static-Valid; Static-I = Static-Invalid
=====

A:Dut-C# show router origin-validation database summary
=====
Static and Dynamic VRP Database Summary
=====
Type                                     IPv4 Routes      IPv6 Routes
-----
192.168.1.1-B                           1                 0
Static                                   1                 0
=====
```

rpki-session

- Syntax** `rpki-session [ip-address] [detail]`
- Context** `show>router>origin-validation`
- Description** This command displays RPKI session information.
- Parameters** *ip-address* — displays RPKI session information for the specified IP address

Values

ipv4-address: a.b.c.d
ipv6-address x:x:x:x:x:x:x
 x:x:x:x:x:d.d.d.d
 where:
 x: [0 to FFFF]H
 d: [0 to 255]D
 interface: 32 chars max, and mandatory for link local
 addresses.

detail — displays the longer, more detailed version of the output

Output The following output is an example of RPKI session information.

Sample Output

```
A:Dut-C# show router origin-validation rpki-session detail
=====
Rpki Session Information
=====
IP Address      : 192.168.1.1
```

```
-----
Port           : 323           Oper State       : established
UpTime         : 0d 00:57:41   Flaps          : 0
Active IPv4 records: 17023     Active IPv6 records: 2515
Admin State    : Up           Local Address   : n/a
Admin State    : Up           Local Address   : 192.0.2.2
Hold Time     : 120           Refresh Time    : 60
Stale Route Time : 3600       Connect Retry   : 120
Serial ID      : 41690        Session ID      : 1452020198
=====
No. of Rpkj-Sessions : 1
=====
```

policy

| | |
|--------------------|--|
| Syntax | policy [<i>name</i> damping prefix-list <i>name</i> as-path <i>name</i> community <i>name</i> admin] |
| Context | show>router |
| Description | This command displays policy-related information. |
| Parameters | name — specifies an existing policy-statement name damping — specifies damping to display route damping profiles prefix-list <i>name</i> — specifies a prefix list name to display the route policy entries as-path <i>name</i> — specifies the route policy AS path name to display route policy entries community <i>name</i> — specifies a route policy community name to display information about a specific community member admin — specifies the admin keyword to display the entities configured in the config>router>policy-options context |
| Output | The following output is an example of router policy information, and Table 25 describes the fields. |

Sample Output

```
B:CORE2# show router policy
=====
Route Policies
=====
Policy           Description
-----
fromStatic
Policies : 1
=====
B:CORE2#
```

Table 25 Policy Fields

| Label | Description |
|-------------|---|
| Policy | The policy name. |
| Description | Displays the description of the policy. |

policy-edits

| | |
|--------------------|--|
| Syntax | policy-edits |
| Context | show>router |
| Description | This command displays edited policy information. |

route-table

| | |
|--------------------|--|
| Syntax | route-table [<i>family</i>] [<i>ip-prefix[/prefix-length]</i>] [longer exact protocol <i>protocol-name</i>] [all] [next-hop-type <i>type</i>] [qos] [alternative] [accounting-class] route-table [<i>family</i>] summary route-table <i>tunnel-endpoints</i> [ip-prefix[/prefix-length]] [longer exact] [detail] route-table [< <i>family</i> >] [< <i>ip-prefix[/prefix-length]</i> >] [longer exact protocol < <i>protocol-name</i> >] extensive [all] |
| Context | show>router |
| Description | This command displays the active routes in the routing table. If no command line arguments are specified, all routes are displayed, sorted by prefix. |
| Parameters | <i>family</i> — specifies the type of routing information to be distributed by this peer group Values ipv4 — displays only those BGP peers that have the IPv4 family enabled and not those capable of exchanging IP-VPN routes ipv6 — displays the BGP peers that are IPv6 capable mcast-ipv4 — displays the BGP peers that are IPv4 multicast capable mcast-ipv6 — displays multicast IPv6 route table <i>ip-prefix[/prefix-length]</i> — displays routes only matching the specified <i>ip-address</i> and length Values The following values apply to the 7750 SR and 7950 XRS: |

```

ipv4-prefix:      a.b.c.d (host bits must be set
                  to 0)
ipv4-prefix-length:      0 to 32
ipv6              ipv6-prefix[/pref*:      x:x:x:x:x:x:x (eight 16-bit
                  pieces)
                  x:x:x:x:x:d.d.d.d
                  x:  [0 to FFFF]H
                  d:  [0 to 255]D
                  prefix-length:      1 to 128ipv6

```

Values The following values apply to the 7450 ESS:

```

ipv4-prefix:      a.b.c.d (host bits must be set to 0)
ipv4-prefix-length:      0 to 32

```

longer — displays routes matching the *ip-prefix/mask* and routes with longer masks

exact — displays the exact route matching the *ip-prefix/mask* masks

protocol *protocol-name* — displays routes learned from the specified protocol

Values local, sub-mgmt, managed, static, ospf, ospf3, isis, rip, aggregate, bgp, bgp-vpn

summary — displays a route table summary information

tunnel-endpoints — specifies to include tunnel endpoint information

next-hop-type tunneled — displays only the tunneled next-hops. For each route entry, the tunnel owner and tunnel ID is shown

Output **Standard Route Table Output** — The following output is an example of standard route table information, and [Table 26](#) describes the fields.

Sample Output

```
*A:Dut-B#config>service>vprn# show router 1 route-table
```

```

=====
Route Table (Service: 1)
=====
Dest Prefix[Flags]                                Type  Proto  Age           Pref
      Next Hop[Interface Name]                      Metric
-----
10.0.0.0/30                                         Local   Local   02h09m23s    0
      to_4007                                         0
10.0.0.8/30                                         Remote  BGP VPN  00h06m38s    170
      1.1.1.9 (tunneled)                             0
11.0.0.8/30                                         Remote  BGP VPN  00h06m38s    170
      1.1.1.9 (tunneled)                             0
192.168.0.0/16 [E]                                 Remote  BGP VPN  00h06m38s    170
      1.1.1.9 (tunneled)                             0

```

```

192.168.0.0/16 [E]                                Remote  BGP VPN    00h06m38s  170
      2.1.1.9 (tunneled)                                0
-----

```

No. of Routes: 4

Flags: L = LFA nexthop available B = BGP backup route available

E = best-external BGP route available

n = Number of times nexthop is repeated

```

*A:Dut-B#config>service>vprn# show router 1 route-table alternative

```

```

=====
Route Table (Service: 1)
=====

```

| Dest Prefix[Flags] | Type | Proto | Age | Pref |
|---------------------------|--------|---------|------------|------|
| Next Hop[Interface Name] | | | Metric | |
| Alt-NextHop | | | Alt-Metric | |
| 10.0.0.0/30 | Local | Local | 02h17m23s | 0 |
| to_4007 | | | 0 | |
| 10.0.0.8/30 | Remote | BGP VPN | 00h14m37s | 170 |
| 1.1.1.9 (tunneled) | | | 0 | |
| 11.0.0.8/30 | Remote | BGP VPN | 00h14m37s | 170 |
| 1.1.1.9 (tunneled) | | | 0 | |
| 192.168.0.0/16 | Remote | BGP VPN | 00h14m37s | 170 |
| 1.1.1.9 (tunneled) | | | 0 | |
| 192.168.0.0/16 (Backup) | Remote | BGP VPN | 00h14m37s | 170 |
| 2.1.1.9 (tunneled) | | | 0 | |
| 192.168.0.0/16 (Best-ext) | Remote | BGP | 00h24m37s | 170 |
| 10.0.0.9 | | | 0 | |

No. of Routes: 5

Flags: Backup = BGP backup route LFA = Loop-Free Alternate nexthop

Best-ext = best-external BGP route

n = Number of times nexthop is repeated

```

*A:Dut-B# show router route-table

```

```

=====
Route Table (Router: Base)
=====

```

| Dest Prefix[Flags] | Type | Proto | Age | Pref |
|--------------------------|--------|-------|-----------|--------|
| Next Hop[Interface Name] | | | | Metric |
| 10.10.1.0/24 | Local | Local | 00h01m25s | 0 |
| ip-10.10.1.2 | | | 0 | |
| 10.10.2.0/24 [L] | Remote | ISIS | 00h00m58s | 15 |
| 10.10.12.3 | | | 13 | |
| 10.10.3.0/24 | Local | Local | 00h01m25s | 0 |
| ip-10.10.3.2 | | | 0 | |
| 10.10.4.0/24 | Local | Local | 00h01m25s | 0 |
| ip-10.10.4.2 | | | 0 | |
| 10.10.5.0/24 [L] | Remote | ISIS | 00h00m58s | 15 |
| 10.10.12.3 | | | 13 | |
| 10.10.6.0/24 [L] | Remote | ISIS | 00h00m58s | 15 |
| 10.10.4.4 | | | 20 | |
| 10.10.9.0/24 [L] | Remote | ISIS | 00h00m58s | 15 |

```

10.10.4.4 20
10.10.10.0/24 [L] Remote ISIS 00h00m58s 15
10.10.12.3 23
10.10.11.0/24 [L] Remote ISIS 00h00m58s 15
10.10.12.3 13
10.10.12.0/24 Local Local 00h01m25s 0
ip-10.10.12.2 0
10.20.1.1/32 [L] Remote ISIS 00h00m58s 15
10.10.1.1 10
10.20.1.2/32 Local Local 00h01m25s 0
system 0
10.20.1.3/32 [L] Remote ISIS 00h00m58s 15
10.10.12.3 3
10.20.1.4/32 [L] Remote ISIS 00h00m58s 15
10.10.4.4 10
10.20.1.5/32 [L] Remote ISIS 00h00m58s 15
10.10.12.3 13
10.20.1.6/32 [L] Remote ISIS 00h00m58s 15
10.10.4.4 20
-----
No. of Routes: 16
Flags: L = LFA nexthop available B = BGP backup route available
=====

*A:Dut-B# show router route-table alternative
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags] Type Proto Age Pref
Next Hop[Interface Name] Metric
Alt-NextHop Alt-Metric
-----
10.10.1.0/24 Local Local 00h02m28s 0
ip-10.10.1.2 0
10.10.2.0/24 Remote ISIS 00h02m01s 15
10.10.12.3 13
10.10.1.1 (LFA) 20
10.10.3.0/24 Local Local 00h02m27s 0
ip-10.10.3.2 0
10.10.4.0/24 Local Local 00h02m28s 0
ip-10.10.4.2 0
10.10.5.0/24 Remote ISIS 00h02m01s 15
10.10.12.3 13
10.10.1.1 (LFA) 20
10.10.6.0/24 Remote ISIS 00h02m01s 15
10.10.4.4 20
10.10.12.3 (LFA) 13
10.10.9.0/24 Remote ISIS 00h02m01s 15
10.10.4.4 20
10.10.12.3 (LFA) 13
10.10.10.0/24 Remote ISIS 00h02m01s 15
10.10.12.3 23
10.10.4.4 (LFA) 20
10.10.11.0/24 Remote ISIS 00h02m01s 15
10.10.12.3 13
10.10.1.1 (LFA) 20
10.10.12.0/24 Local Local 00h02m28s 0
ip-10.10.12.2 0
10.20.1.1/32 Remote ISIS 00h02m01s 15

```

```
10.10.1.1 10
10.10.12.3 (LFA) 13
10.20.1.2/32 Local Local 00h02m28s 0
system 0
10.20.1.3/32 Remote ISIS 00h02m05s 15
10.10.12.3 3
10.10.1.1 (LFA) 20
10.20.1.4/32 Remote ISIS 00h02m05s 15
10.10.4.4 10
10.10.12.3 (LFA) 13
10.20.1.5/32 Remote ISIS 00h02m05s 15
10.10.12.3 13
10.10.4.4 (LFA) 20
10.20.1.6/32 Remote ISIS 00h02m05s 15
10.10.4.4 20
10.10.12.3 (LFA) 23
-----
No. of Routes: 16
Flags: Backup = BGP backup routeLFA = Loop-Free Alternate nexthop
=====
```

```
*A:Dut-C# show router route-table 1.1.1.1/32
=====
Route Table (Router: Base)
=====
Dest Prefix                                Type  Proto  Age           Pref
      Next Hop[Interface Name]                                Metric
-----
1.1.1.1/32                                Remote BGP    00h00m09s    170
      10.20.1.1 (tunneled:RSVP:1)                                0
-----
No. of Routes: 1
=====
```

```
A:ALA# show router route-table
=====
Route Table (Router: Base)
=====
Dest Prefix                                Type  Proto
Age      Pref
      Next Hop[Interface Name]                                Metric
-----
11.2.103.0/24                                Remote OSPF
00h59m02s  10
      21.2.4.2                                2
11.2.103.0/24                                Remote OSPF
00h59m02s  10
      22.2.4.2                                2
11.2.103.0/24                                Remote OSPF
00h59m02s  10
      23.2.4.2                                2
11.2.103.0/24                                Remote OSPF
00h59m02s  10
      24.2.4.2                                2
11.2.103.0/24                                Remote OSPF
00h59m02s  10
      100.0.0.1                                2
11.2.103.0/24                                Remote OSPF
00h59m02s  10
```



```

100.128.0.1
11.4.101.0/24 Local Local 02h14m29s 0
...
-----
A:ALA#

B:ALA-B# show router route-table 100.10.0.0 exact
=====
Dest Address Next Hop Type Proto Age Metric Pref
-----
100.10.0.0/16 Black Hole Remote Static 00h03m17s 1 5
-----
No. of Routes: 1
=====

B:ALA-B#
A:ALA-A# show router route-table 10.10.0.4
=====
Route Table
=====
Dest Address Next Hop Type Protocol Age Metric Pref
-----
10.10.0.4/32 10.10.34.4 Remote OSPF 3523 1001 10
-----
A:ALA-A#

A:ALA-A# show router route-table 10.10.0.4/32 longer
=====
Route Table
=====
Dest Address Next Hop Type Protocol Age Metric Pref
-----
10.10.0.4/32 10.10.34.4 Remote OSPF 3523 1001 10
-----
No. of Routes: 1
=====
+ : indicates that the route matches on a longer prefix
A:ALA-A#

*A:Dut-C# show router route-table
=====
Route Table (Router: Base)
=====
Dest Prefix[Flags] Type Proto Age Pref
Next Hop[Interface Name] Metric
-----
1.1.2.0/24 Remote ISIS 00h44m24s 15
1.1.3.1 20
1.1.2.0/24 Remote ISIS 00h44m24s 15
1.2.3.2 20
1.1.3.0/24 Local Local 00h44m30s 0
to_Dut-A 0
1.1.9.0/24 Remote ISIS 00h44m16s 15
1.1.3.1 20
1.2.3.0/24 Local Local 00h44m30s 0
to_Dut-B 0
1.2.9.0/24 Remote ISIS 00h43m55s 160
1.2.3.2 10
10.12.0.0/24 Local Local 00h44m29s 0

```

```

        itfToArborCP_02                                0
10.20.1.1/32                                           Remote  ISIS    00h44m24s  15
        1.1.3.1                                         10
10.20.1.2/32                                           Remote  ISIS    00h44m28s  15
        1.2.3.2                                         10
10.20.1.3/32                                           Local   Local   00h44m32s   0
        system                                           0
20.12.0.43/32                                           Remote  Static  00h44m31s   5
        vprn1:mda-1-1                                   1
20.12.0.44/32                                           Remote  Static  00h44m31s   5
        vprn1:mda-2-1                                   1
20.12.0.45/32                                           Remote  Static  00h44m31s   5
        vprn1:mda-2-2                                   1
20.12.0.46/32                                           Remote  Static  00h44m30s   5
        vprn1:mda-3-1                                   1
138.203.71.202/32                                       Remote  Static  00h44m29s   5
        10.12.0.2                                       1

```

No. of Routes: 17

Flags: L = LFA nexthop available B = BGP backup route available

n = Number of times nexthop is repeated

=====

A:ALA-A# show router route-table protocol ospf

=====

Route Table

```

=====
Dest Address      Next Hop      Type   Protocol   Age      Metric  Pref
-----
10.10.0.1/32      10.10.13.1    Remote OSPF       65844    1001    10
10.10.0.2/32      10.10.13.1    Remote OSPF       65844    2001    10
10.10.0.4/32      10.10.34.4    Remote OSPF       3523     1001    10
10.10.0.5/32      10.10.35.5    Remote OSPF    1084022    1001    10
10.10.12.0/24     10.10.13.1    Remote OSPF       65844    2000    10
10.10.15.0/24     10.10.13.1    Remote OSPF       58836    2000    10
10.10.24.0/24     10.10.34.4    Remote OSPF       3523     2000    10
10.10.25.0/24     10.10.35.5    Remote OSPF    399059    2000    10
10.10.45.0/24     10.10.34.4    Remote OSPF       3523     2000    10
=====

```

A:ALA-A#

show router route-table 131.132.133.134/32 next-hop-type tunneled

Route Table (Router: Base)

```

=====
Dest Prefix      Type      Proto   Age      Pre
f
        Next Hop[Interface Name]      Metric
-----
131.132.133.134/32      Remote  OSPF    00h02m09s  10
        66.66.66.66      10
        Next-hop type: tunneled, Owner: RSVP, Tunnel-ID: <out-ifindex-from-route>
=====

```

No. of Routes: 1

=====

*A:Dut-B# show router route-table next-hop-type tunneled

=====

Route Table (Router: Base)

```

=====
Dest Prefix      Type      Proto   Age      Pref

```

| Next Hop[Interface Name] | | Metric | | |
|-----------------------------|--------|--------|-----------|----|
| 10.10.5.0/24 | Remote | OSPF | 00h02m20s | 10 |
| 10.20.1.5 (tunneled:RSVP:1) | | | 1100 | |
| 10.10.10.0/24 | Remote | OSPF | 00h02m20s | 10 |
| 10.20.1.5 (tunneled:RSVP:1) | | | 1100 | |
| 10.20.1.5/32 | Remote | OSPF | 00h02m20s | 10 |
| 10.20.1.5 (tunneled:RSVP:1) | | | 100 | |
| 10.20.1.6/32 | Remote | OSPF | 00h02m20s | 10 |
| 10.20.1.5 (tunneled:RSVP:1) | | | 1100 | |

No. of Routes: 4

*A:Dut-B# show router route-table 10.20.1.5/32 next-hop-type tunneled

Route Table (Router: Base)

| Dest Prefix | Type | Proto | Age | Pref |
|-----------------------------|--------|--------|-----------|------|
| Next Hop[Interface Name] | | Metric | | |
| 10.20.1.5/32 | Remote | OSPF | 00h03m55s | 10 |
| 10.20.1.5 (tunneled:RSVP:1) | | | 100 | |

No. of Routes: 1

*A:Dut-C#

*A:Dut-C# show router route-table summary

Route Table Summary (Router: Base)

| | Active | Available |
|--------------|--------|-----------|
| Static | 5 | 5 |
| Direct | 12 | 12 |
| Host | 0 | 11 |
| BGP | 0 | 0 |
| BGP (Backup) | 0 | 0 |
| VPN Leak | 0 | 0 |
| OSPF | 0 | 0 |
| ISIS | 6 | 6 |
| ISIS (LFA) | 0 | 0 |
| RIP | 0 | 0 |
| LDP | 0 | 0 |
| Aggregate | 0 | 0 |
| Sub Mgmt | 0 | 0 |
| Managed | 0 | 0 |
| NAT | 0 | 0 |
| Total | 24 | 35 |

NOTE: ISIS LFA routes and BGP Backup routes are not counted towards the total.

*A:Dut-C# show router route-table ipv6 3ffe::10:20:1:4/128

IPv6 Route Table (Router: Base)

| Dest Prefix[Flags] | Type | Proto | Age | Pref |
|--------------------|------|-------|-----|------|
|--------------------|------|-------|-----|------|

```

      Next Hop[Interface Name]
-----
3ffe::10:20:1:4/128          Remote  ISIS    00h12m48s  15
      fe80::205e:1ff:fe01:1-"ip-10.10.5.3"
                               20
3ffe::10:20:1:4/128          Remote  ISIS    00h12m48s  15
      fe80::6629:ffff:fe00:141-"ip-10.10.12.3"
                               20
-----

```

No. of Routes: 2

Flags: n = Number of times nexthop is repeated

B = BGP backup route available

L = LFA nexthop available

S = Sticky ECMP requested

```
=====
*A:Dut-C>config>router>static-route-entry#

```

```
*A:Dut-C# show router route-table ipv6 3ffe::10:20:1:4/128 extensive

```

```
=====
Route Table (Router: Base)
=====

```

```

Dest Prefix      : 3ffe::10:20:1:4/128
Protocol         : ISIS
Age              : 00h12m55s
Preference       : 15
Next-Hop         : fe80::205e:1ff:fe01:1-"ip-10.10.5.3"
  Interface      : ip-10.10.5.3
  QoS            : Priority=n/c, FC=n/c
  Source-Class   : 0
  Dest-Class     : 0
  Metric         : 20
  ECMP-Weight    : N/A
Next-Hop         : fe80::6629:ffff:fe00:141-"ip-10.10.12.3"
  Interface      : ip-10.10.12.3
  QoS            : Priority=n/c, FC=n/c
  Source-Class   : 0
  Dest-Class     : 0
  Metric         : 20
  ECMP-Weight    : N/A
-----

```

No. of Destinations: 1

```
=====
*A:Dut-C# show router route-table ipv6 3ffe::10:20:1:4/128 all

```

```
=====
IPv6 Route Table (Router: Base)
=====

```

```

Dest Prefix[Flags]      Type  Proto  Age      Pref
      Next Hop[Interface Name]      Active  Metric
-----
3ffe::10:20:1:4/128     Remote ISIS    00h13m00s  15
      fe80::205e:1ff:fe01:1-"ip-10.10.5.3"
                               Y        20
3ffe::10:20:1:4/128     Remote ISIS    00h13m00s  15
      fe80::6629:ffff:fe00:141-"ip-10.10.12.3"
                               Y        20
3ffe::10:20:1:4/128     Remote ISIS(1) 00h13m09s  15
      fe80::205e:1ff:fe01:1-"ip-10.10.5.3"
                               N        20
3ffe::10:20:1:4/128     Remote ISIS(1) 00h13m09s  15
      fe80::6629:ffff:fe00:141-"ip-10.10.12.3"
                               N        20
-----

```

No. of Routes: 4

```

Flags: n = Number of times nexthop is repeated
      B = BGP backup route available
      L = LFA nexthop available
      S = Sticky ECMP requested
      E = Inactive best-external BGP route
=====

*A:Dut-C# show router route-table ipv6 3ffe::10:20:1:4/128 all extensive
=====
Route Table (Router: Base)
=====
Dest Prefix      : 3ffe::10:20:1:4/128
Protocol         : ISIS
Age              : 00h13m06s
Preference       : 15
Next-Hop         : fe80::205e:1ff:fe01:1-"ip-10.10.5.3"
  Interface      : ip-10.10.5.3
  Active         : Yes
  QoS             : Priority=n/c, FC=n/c
  Source-Class    : 0
  Dest-Class      : 0
  Metric          : 20
  ECMP-Weight     : N/A
Next-Hop         : fe80::6629:ffff:fe00:141-"ip-10.10.12.3"
  Interface      : ip-10.10.12.3
  Active         : Yes
  QoS             : Priority=n/c, FC=n/c
  Source-Class    : 0
  Dest-Class      : 0
  Metric          : 20
  ECMP-Weight     : N/A
-----
Dest Prefix      : 3ffe::10:20:1:4/128
Protocol         : ISIS (1)
Age              : 00h13m15s
Preference       : 15
Next-Hop         : fe80::205e:1ff:fe01:1-"ip-10.10.5.3"
  Interface      : ip-10.10.5.3
  Active         : No
  QoS             : Priority=n/c, FC=n/c
  Source-Class    : 0
  Dest-Class      : 0
  Metric          : 20
  ECMP-Weight     : N/A
Next-Hop         : fe80::6629:ffff:fe00:141-"ip-10.10.12.3"
  Interface      : ip-10.10.12.3
  Active         : No
  QoS             : Priority=n/c, FC=n/c
  Source-Class    : 0
  Dest-Class      : 0
  Metric          : 20
  ECMP-Weight     : N/A
-----
No. of Destinations: 2
=====

```

Table 26 Standard Route Table Fields

| Label | Description |
|---------------|--|
| Dest Address | The route destination address and mask. |
| Next Hop | The next hop IP address for the route destination. |
| Type | Local The route is a local route. |
| | Remote The route is a remote route. |
| Protocol | The protocol through which the route was learned. |
| Age | The route age in seconds for the route. |
| Metric | The route metric value for the route. |
| Pref | The route preference value for the route. |
| No. of Routes | The number of routes displayed in the list. |

Summary Route Table Output — Summary output for the route table displays the number of active routes and the number of routes learned by the router by protocol. Total active and available routes are also displayed.

The following output is an example of summary route table information.

Sample Output

```
A:ALA-A# show router route-table summary
=====
Route Table Summary
=====

```

| | Active | Available |
|-----------|--------|-----------|
| Static | 1 | 1 |
| Direct | 6 | 6 |
| BGP | 0 | 0 |
| OSPF | 9 | 9 |
| ISIS | 0 | 0 |
| RIP | 0 | 0 |
| Aggregate | 0 | 0 |
| Total | 16 | 16 |

```
=====
A:ALA-A#

*A:SRR# show router route-table summary
=====
Route Table Summary (Router: Base)
=====
```

| | Active | Available |
|--------------|--------|-----------|
| Static | 6 | 6 |
| Direct | 1698 | 1698 |
| Host | 0 | 1477 |
| BGP | 0 | 0 |
| BGP (Backup) | 0 | 0 |
| VPN Leak | 0 | 0 |
| OSPF | 0 | 0 |
| ISIS | 3296 | 6383 |
| ISIS (LFA) | 472 | 1499 |
| RIP | 0 | 0 |
| LDP | 6 | 6 |
| Aggregate | 0 | 0 |
| Sub Mgmt | 0 | 0 |
| Managed | 0 | 0 |
| NAT | 0 | 0 |
| Total | 5006 | 9570 |

NOTE: ISIS LFA routes and BGP Backup routes are not counted towards the total.
*A:SRR#
*A:Dut-C>config>router>mpls>lsp# show router route-table 10.0.0.2/32 extensive

Route Table (Router: Base)

```

Dest Prefix      : 10.0.0.2/32
Protocol         : OSPF (1)
Age              : 00h02m40s
Preference      : 150
Next-Hop         : 1.0.0.3 (RSVP tunnel:94)
  QoS             : Priority=n/c, FC=n/c
  Source-Class    : 0
  Dest-Class      : 0
  Metric          : 10
  ECMP-Weight     : 20
Next-Hop         : 1.0.0.3 (RSVP tunnel:61442)
  QoS             : Priority=n/c, FC=n/c
  Source-Class    : 0
  Dest-Class      : 0
  Metric          : 10
  ECMP-Weight     : 1

```

No. of Destinations: 1

*A:Dut-C>config>router>static-route-entry>indirect>tunnel-next-hop# show router route-table 10.1.0.5/32 extensive

Route Table (Router: Base)

```

Dest Prefix      : 10.1.0.5/32
Protocol         : STATIC
Age              : 00h00m11s
Preference      : 5
Next-Hop         : 1.0.0.2 (RSVP tunnel:128)
  QoS             : Priority=n/c, FC=n/c
  Source-Class    : 0

```

```

      Dest-Class      : 0
      Metric          : 1
      ECMP-Weight     : 18
Next-Hop             : 1.0.0.2 (RSVP tunnel:132)
      QoS             : Priority=n/c, FC=n/c
      Source-Class    : 0
      Dest-Class      : 0
      Metric          : 1
      ECMP-Weight     : 2
Next-Hop             : 1.0.0.3 (RSVP tunnel:94)
      QoS             : Priority=n/c, FC=n/c
      Source-Class    : 0
      Dest-Class      : 0
      Metric          : 1
      ECMP-Weight     : 7
Next-Hop             : 1.0.0.3 (RSVP tunnel:61442)
      QoS             : Priority=n/c, FC=n/c
      Source-Class    : 0
      Dest-Class      : 0
      Metric          : 1
      ECMP-Weight     : 2
-----
No. of Destinations: 1
=====
```

rtr-advertisement

| | | | |
|-------------|---|---------------------|---|
| Syntax | rtr-advertisement [interface <i>interface-name</i>] [prefix <i>ipv6-prefix[/prefix-length]</i>] rtr-advertisement [conflicts] | | |
| Context | show>router | | |
| Description | This command displays router advertisement information. If no command line arguments are specified, all routes are displayed, sorted by prefix. | | |
| Parameters | <i>interface-name</i> — maximum 32 characters <i>ipv6-prefix[/prefix-length]</i> — displays routes only matching the specified <i>ip-address</i> and length and only applies to the 7750 SR and 7950 XRS | | |
| | Values | | |
| | ipv6 | ipv6-prefix[/pref*: | x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0 to FFFF]H d: [0 to 255]D |
| | | prefix-length: | 1 to 128 |
| Output | Router-Advertisement Table Output — The following output is an example of router advertisement information, and Table 27 describes the fields. | | |

Sample Output

```
A:Dut-A# show router rtr-advertisement
=====
Router Advertisement
=====
-----
Interface: interfaceNetworkNonDefault
-----
Rtr Advertisement Tx : 8                Last Sent           : 00h01m28s
Nbr Solicitation Tx  : 83               Last Sent           : 00h00m17s
Nbr Advertisement Tx : 74               Last Sent           : 00h00m25s
Rtr Advertisement Rx : 8                Rtr Solicitation Rx : 0
Nbr Advertisement Rx : 83               Nbr Solicitation Rx : 74
-----
Server1               : 2001:db8::1
Server2               : N/A
Server3               : N/A
Server4               : N/A
Rdnss-lifetime        : 1200            Include-dns         : yes
-----
Max Advert Interval   : 601             Min Advert Interval : 201
Managed Config       : TRUE             Other Config        : TRUE
Reachable Time        : 00h00m00s400ms   Router Lifetime     : 00h30m01s
Retransmit Time       : 00h00m00s400ms   Hop Limit           : 63
Link MTU              : 1500
-----
Prefix: 211::/120
Autonomous Flag       : FALSE            On-link flag        : FALSE
Preferred Lifetime    : 07d00h00m        Valid Lifetime      : 30d00h00m
-----
Prefix: 231::/120
Autonomous Flag       : FALSE            On-link flag        : FALSE
Preferred Lifetime    : 49710d06h        Valid Lifetime      : 49710d06h
-----
Prefix: 241::/120
Autonomous Flag       : TRUE             On-link flag        : TRUE
Preferred Lifetime    : 00h00m00s        Valid Lifetime      : 00h00m00s
-----
Prefix: 251::/120
Autonomous Flag       : TRUE             On-link flag        : TRUE
Preferred Lifetime    : 07d00h00m        Valid Lifetime      : 30d00h00m
-----
Advertisement from: FE80::200:FF:FE00:2
Managed Config       : FALSE            Other Config        : FALSE
Reachable Time        : 00h00m00s0ms     Router Lifetime     : 00h30m00s
Retransmit Time       : 00h00m00s0ms     Hop Limit           : 64
Link MTU              : 0
-----
Interface: interfaceServiceNonDefault
-----
Rtr Advertisement Tx : 8                Last Sent           : 00h06m41s
Nbr Solicitation Tx  : 166               Last Sent           : 00h00m04s
Nbr Advertisement Tx : 143               Last Sent           : 00h00m05s
Rtr Advertisement Rx : 8                Rtr Solicitation Rx : 0
Nbr Advertisement Rx : 166               Nbr Solicitation Rx : 143
-----
Max Advert Interval   : 601             Min Advert Interval : 201
Managed Config       : TRUE             Other Config        : TRUE
```

```

Reachable Time      : 00h00m00s400ms  Router Lifetime     : 00h30m01s
Retransmit Time     : 00h00m00s400ms  Hop Limit          : 63
Link MTU            : 1500

Prefix: 23::/120
Autonomous Flag     : FALSE             On-link flag        : FALSE
Preferred Lifetime  : infinite          Valid Lifetime      : infinite

Prefix: 24::/120
Autonomous Flag     : TRUE              On-link flag        : TRUE
Preferred Lifetime  : 00h00m00s         Valid Lifetime      : 00h00m00s

Prefix: 25::/120
Autonomous Flag     : TRUE              On-link flag        : TRUE
Preferred Lifetime  : 07d00h00m         Valid Lifetime      : 30d00h00m
-----
Advertisement from: FE80::200:FF:FE00:2
Managed Config     : FALSE             Other Config        : FALSE
Reachable Time     : 00h00m00s0ms       Router Lifetime     : 00h30m00s
Retransmit Time     : 00h00m00s0ms       Hop Limit          : 64
Link MTU           : 0

Prefix: 2::/120
Autonomous Flag     : TRUE             On-link flag        : TRUE
Preferred Lifetime  : 07d00h00m         Valid Lifetime      : 30d00h00m

Prefix: 23::/120
Autonomous Flag     : TRUE             On-link flag        : TRUE
Preferred Lifetime  : 07d00h00m         Valid Lifetime      : 30d00h00m

Prefix: 24::/119
Autonomous Flag     : TRUE             On-link flag        : TRUE
Preferred Lifetime  : 07d00h00m         Valid Lifetime      : 30d00h00m

Prefix: 25::/120
Autonomous Flag     : TRUE             On-link flag        : TRUE
Preferred Lifetime  : 07d00h00m         Valid Lifetime      : infinite

Prefix: 231::/120
Autonomous Flag     : TRUE             On-link flag        : TRUE
Preferred Lifetime  : 07d00h00m         Valid Lifetime      : 30d00h00m
-----
...
A:Dut-A#

```

Table 27 Router Advertisement Table Fields

| Label | Description |
|------------------------------------|--|
| Rtr Advertisement Tx/ Last Sent | The number of router advertisements sent and time since they were sent. |
| Nbr Solicitation Tx | The number of neighbor solicitations sent and time since they were sent. |

Table 27 Router Advertisement Table Fields (Continued)

| Label | Description |
|----------------------|---|
| Nbr Advertisement Tx | The number of neighbor advertisements sent and time since they were sent. |
| Rtr Advertisement Rx | The number of router advertisements received and time since they were received. |
| Nbr Advertisement Rx | The number of neighbor advertisements received and time since they were received. |
| Max Advert Interval | The maximum interval between sending router advertisement messages. |
| Managed Config | True Indicates that DHCPv6 has been configured. |
| | False Indicates that DHCPv6 is not available for address configuration. |
| Reachable Time | The time, in milliseconds, that a node assumes a neighbor is reachable after receiving a reachability confirmation. |
| Retransmit Time | The time, in milliseconds, between retransmitted neighbor solicitation messages. |
| Link MTU | The MTU number the nodes use for sending packets on the link. |
| Rtr Solicitation Rx | The number of router solicitations received and time since they were received. |
| Nbr Solicitation Rx | The number of neighbor solicitations received and time since they were received. |
| Min Advert Interval | The minimum interval between sending ICMPv6 neighbor discovery router advertisement messages. |
| Other Config | True Indicates there are other stateful configurations. |
| | False Indicates there are no other stateful configurations. |
| Router Lifetime | Displays the router lifetime in seconds. |
| Hop Limit | Displays the current hop limit. |

Router-Advertisement Conflicts Output — The following output is an example of router advertisement conflicts, and [Table 28](#) describes the fields.

Sample Output

```

A:Dut-A# show>router# rtr-advertisement conflicts
=====
Router Advertisement
=====
Interface: interfaceNetworkNonDefault
-----
Advertisement from: FE80::200:FF:FE00:2
Managed Config      : FALSE [TRUE]
Other Config         : FALSE [TRUE]
Reachable Time       : 00h00m00s0ms [00h00m00s400ms]
Router Lifetime      : 00h30m00s [00h30m01s]
Retransmit Time      : 00h00m00s0ms [00h00m00s400ms]
Hop Limit            : 64 [63]
Link MTU             : 0 [1500]

Prefix not present in neighbor router advertisement
Prefix: 211::/120
Autonomous Flag      : FALSE          On-link flag      : FALSE
Preferred Lifetime   : 07d00h00m      Valid Lifetime    : 30d00h00m

Prefix not present in neighbor router advertisement
Prefix: 231::/120
Autonomous Flag      : FALSE          On-link flag      : FALSE
Preferred Lifetime   : 49710d06h      Valid Lifetime    : 49710d06h

Prefix not present in neighbor router advertisement
Prefix: 241::/120
Autonomous Flag      : TRUE           On-link flag      : TRUE
Preferred Lifetime   : 00h00m00s      Valid Lifetime    : 00h00m00s

Prefix not present in neighbor router advertisement
Prefix: 251::/120
Autonomous Flag      : TRUE           On-link flag      : TRUE
Preferred Lifetime   : 07d00h00m      Valid Lifetime    : 30d00h00m
-----
Interface: interfaceServiceNonDefault
-----
Advertisement from: FE80::200:FF:FE00:2
Managed Config      : FALSE [TRUE]
Other Config         : FALSE [TRUE]
Reachable Time       : 00h00m00s0ms [00h00m00s400ms]
Router Lifetime      : 00h30m00s [00h30m01s]
Retransmit Time      : 00h00m00s0ms [00h00m00s400ms]
Hop Limit            : 64 [63]
Link MTU             : 0 [1500]

Prefix not present in own router advertisement
Prefix: 2::/120
Autonomous Flag      : TRUE           On-link flag      : TRUE
Preferred Lifetime   : 07d00h00m      Valid Lifetime    : 30d00h00m

Prefix: 23::/120
Autonomous Flag      : TRUE [FALSE]
On-link flag         : TRUE [FALSE]
Preferred Lifetime   : 07d00h00m [infinite]
Valid Lifetime       : 30d00h00m [infinite]

```

```

Prefix not present in own router advertisement
Prefix: 24::/119
Autonomous Flag      : TRUE          On-link flag      : TRUE
Preferred Lifetime   : 07d00h00m     Valid Lifetime    : 30d00h00m

Prefix not present in neighbor router advertisement
Prefix: 24::/120
Autonomous Flag      : TRUE          On-link flag      : TRUE
Preferred Lifetime   : 00h00m00s     Valid Lifetime    : 00h00m00s

Prefix: 25::/120
Valid Lifetime      : infinite [30d00h00m]

Prefix not present in own router advertisement
Prefix: 231::/120
Autonomous Flag      : TRUE          On-link flag      : TRUE
Preferred Lifetime   : 07d00h00m     Valid Lifetime    : 30d00h00m
=====
A:Dut-A#

```

Table 28 Router-Advertisement Conflicts Fields

| Label | Description |
|--------------------|---|
| Advertisement from | The address of the advertising router. |
| Reachable Time | The time, in milliseconds, that a node assumes a neighbor is reachable after receiving a reachability confirmation. |
| Router Lifetime | Displays the router lifetime in seconds. |
| Retransmit Time | The time, in milliseconds, between retransmitted neighbor solicitation messages. |
| Hop Limit | Displays the current hop limit |
| Link MTU | The MTU number the nodes use for sending packets on the link. |

static-arp

- Syntax** **static-arp** [*ip-addr* | *ip-int-name* | **mac** *ieee-mac-addr*]
- Context** show>router
- Description** This command displays the router static ARP table sorted by IP address. If no options are present, all ARP entries are displayed.
- Parameters**
- ip-addr* — only displays static ARP entries associated with the specified IP address
 - ip-int-name* — only displays static ARP entries associated with the specified IP interface name
 - mac** *ieee-mac-addr* — only displays static ARP entries associated with the specified MAC address

Output **Static ARP Table Output** — The following output is an example of static AARP table information, and [Table 29](#) describes the output fields.

Sample Output

```
A:ALA-A# show router static-arp
=====
ARP Table
=====
IP Address      MAC Address      Age      Type Interface
-----
10.200.0.253    00:00:5a:40:00:01 00:00:00 Sta  to-ser1
12.200.1.1      00:00:5a:01:00:33 00:00:00 Inv  to-ser1a
-----
No. of ARP Entries: 1
=====
A:ALA-A#

A:ALA-A# show router static-arp 12.200.1.1
=====
ARP Table
=====
IP Address      MAC Address      Age      Type Interface
-----
12.200.1.1      00:00:5a:01:00:33 00:00:00 Inv  to-ser1
-----
=====
A:ALA-A#

A:ALA-A# show router static-arp to-ser1
=====
ARP Table
=====
IP Address      MAC Address      Age      Type Interface
-----
10.200.0.253    00:00:5a:40:00:01 00:00:00 Sta  to-
ser1
=====
A:ALA-A#

A:ALA-A# show router static-arp mac 00:00:5a:40:00:01
=====
ARP Table
=====
IP Address      MAC Address      Age      Type Interface
-----
10.200.0.253    00:00:5a:40:00:01 00:00:00 Sta  to-
ser1
=====
A:ALA-A#
```

Table 29 Static ARP Table Fields

| Label | Description |
|--------------------|---|
| IP Address | The IP address of the static ARP entry. |
| MAC Address | The MAC address of the static ARP entry. |
| Age | The age of the ARP entry. Static ARPs always have 00:00:00 for the age. |
| Type | Inv The ARP entry is an inactive static ARP entry (invalid). |
| | Sta The ARP entry is an active static ARP entry. |
| Interface | The IP interface name associated with the ARP entry. |
| No. of ARP Entries | The number of ARP entries displayed in the list. |

static-route

| | |
|--------------------|--|
| Syntax | static-route [family] [[<i>ip-prefix</i> / <i>mask</i>] [preference <i>preference</i>] [next-hop <i>ip-address</i>] tag tag] |
| Context | show>router |
| Description | This command displays the static entries in the routing table. If no options are present, all static routes are displayed sorted by prefix. |
| Parameters | <p>family — specifies the type of routing information to be distributed by this peer group</p> <p>Values</p> <ul style="list-style-type: none"> ipv4 — displays only those BGP peers that have the IPv4 family enabled and not those capable of exchanging IP-VPN routes ipv6 — displays the BGP peers that are IPv6 capable mcast-ipv4 — displays the BGP peers that are IPv4 multicast capable <p><i>ip-prefix</i> /<i>mask</i> — displays static routes only matching the specified <i>ip-prefix</i> and <i>mask</i></p> <p>Values The following values apply to the 7750 SR and 7950 XRS:</p> |

```

ipv4-prefix:          a.b.c.d (host bits must be 0)
                       ipv4-prefix-length:          0 to 32
ipv6-prefix:          x:x:x:x:x:x:x (eight 16-bit pieces)
                       x:x:x:x:x:x:d.d.d.d
                       x:                                [0 to FFFF]H
                       d:                                [0 to 255]D
                       ipv6-prefix-length:          0 to 128

```

Values The following values apply to the 7450 ESS:

```

ipv4-prefix:          a.b.c.d (host bits must be 0)
ipv4-prefix-length:    0 to 32

```

preference *preference* — only displays static routes with the specified route preference

Values 0 to 65535

next-hop *ip-address* — only displays static routes with the specified next hop IP address

Values The following values apply to the 7750 SR and 7950 XRS:

```

ipv4-address:          a.b.c.d (host bits must be 0)
ipv6-address:          x:x:x:x:x:x:x (eight 16-bit
                       pieces)
                       x:x:x:x:x:x:d.d.d.d
                       x:                                [0 to FFFF]H
                       d:                                [0 to 255]D

```

Values The following values apply to the 7450 ESS:

```

ipv4-address: a.b.c.d (host bits must be 0)

```

tag *tag* — displays the tag used to add a 32-bit integer tag to the static route. The tag is used in route policies to control distribution of the route into other protocols.

Values 1 to 4294967295

Output **Static Route Output** — The following output is an example of static route information, and [Table 30](#) describes the fields.

Sample Output

```

A:ALA-A# show router static-route
=====
Route Table
=====
IP Addr/mask      Pref Metric Type Nexthop      Interface      Active
-----
192.168.250.0/24  5      1      ID    10.200.10.1    to-ser1       Y

```



```

192.168.252.0/24  5    1    NH    10.10.0.254      n/a      N
192.168.253.0/24  5    1    NH    to-ser1          n/a      N
192.168.253.0/24  5    1    NH    10.10.0.254      n/a      N
192.168.254.0/24  4    1    BH    black-hole       n/a      Y
=====

```

A:ALA-A#

A:ALA-A# show router static-route 192.168.250.0/24

```

=====
Route Table
=====
IP Addr/mask      Pref Metric Type Nexthop      Interface      Active
-----
192.168.250.0/24  5    1    ID    10.200.10.1      to-ser1        Y
=====

```

A:ALA-A#

A:ALA-A# show router static-route preference 4

```

=====
Route Table
=====
IP Addr/mask      Pref Metric Type Nexthop      Interface      Active
-----
192.168.254.0/24  4    1    BH    black-hole       n/a            Y
=====

```

A:ALA-A#

A:ALA-A# show router static-route next-hop 10.10.0.254

```

=====
Route Table
=====
IP Addr/mask      Pref Metric Type Nexthop      Interface      Active
-----
192.168.253.0/24  5    1    NH    10.10.0.254      n/a            N
=====

```

A:ALA-A#

Static Route Table (Router: Base) Family: IPv6

```

=====
Prefix           : 3ffe::10:10:14:0/120
Nexthop          : 3ffe::10:20:1:6
Type             : Indirect
Interface        : n/a
Active           : Y
Prefix List      : n/a
Prefix List Type : n/a
Metric           : 1
Preference       : 5
Source Class     : 0
Dest Class       : 0
Admin State      : Up
Tag              : 0
Creation Origin  : manual
BFD              : disabled
Community        :
CPE-check        : disabled
Tunnel Resolution: any
Disallow-IGP     : disabled
RSVP-TE Tunnels : disabled
LDP Tunnels      : disabled
SR-ISIS Tunnels : disabled
SR-OSPF Tunnels  : disabled
=====

```

```

SR-TE Tunnels      : disabled
-----
Prefix             : 3ffe::10:10:14:0/120
Nextthop           : 3ffe::20:20:1:6
Type               : Indirect
Interface          : n/a
Prefix List        : n/a
Metric            : 1
Source Class       : 0
Admin State        : Up
Creation Origin    : manual
BFD                : disabled
Community          :
CPE-check          : disabled
Tunnel Resolution  : any
RSVP-TE Tunnels    : disabled
SR-ISIS Tunnels    : disabled
SR-TE Tunnels      : disabled
-----
Disallow-IGP       : disabled
LDP Tunnels        : disabled
SR-OSPF Tunnels    : disabled
-----
Prefix             : 3ffe::10:10:14:0/120
Nextthop           : 3ffe::20:20:1:5
Type               : Indirect
Interface          : n/a
Prefix List        : n/a
Metric            : 1
Source Class       : 0
Admin State        : Up
Creation Origin    : manual
BFD                : disabled
Community          :
CPE-check          : disabled
Tunnel Resolution  : any
RSVP-TE Tunnels    : disabled
SR-ISIS Tunnels    : disabled
SR-TE Tunnels      : disabled
-----
Disallow-IGP       : disabled
LDP Tunnels        : disabled
SR-OSPF Tunnels    : disabled
-----
Prefix             : 3ffe::10:10:14:0/120
Nextthop           : 3ffe::10:20:1:5
Type               : Indirect
Interface          : n/a
Prefix List        : n/a
Metric            : 1
Source Class       : 0
Admin State        : Up
Creation Origin    : manual
BFD                : disabled
Community          :
CPE-check          : disabled
Tunnel Resolution  : filter
RSVP-TE Tunnels    : disabled
SR-ISIS Tunnels    : enabled
SR-TE Tunnels      : disabled
-----
Disallow-IGP       : disabled
LDP Tunnels        : disabled
SR-OSPF Tunnels    : disabled
-----
No. of Static Routes: 4
=====
*A:Dut-C#

```

Table 30 Static Route Fields

| Label | Description |
|---------------|--|
| IP Addr/mask | The static route destination address and mask. |
| Pref | The route preference value for the static route. |
| Metric | The route metric value for the static route. |
| Type | BH The static route is a black hole route. The nexthop for this type of route is black-hole. |
| | ID The static route is an indirect route, where the nexthop for this type of route is the non-directly connected next hop. |
| | NH The route is a static route with a directly connected next hop. The Nexthop for this type of route is either the next hop IP address or an egress IP interface name. |
| Next Hop | The next hop for the static route destination. |
| Protocol | The protocol through which the route was learned. |
| Interface | The egress IP interface name for the static route. n/a indicates there is no current egress interface because the static route is inactive or a black hole route. |
| Active | N The static route is inactive; for example, the static route is disabled or the next hop IP interface is down. |
| | Y The static route is active. |
| No. of Routes | The number of routes displayed in the list. |

The following output is an example static route information for the 7750 SR and 7950 XRS:

```
*A:sim1# show router static-route 10.10.0.0/16 detail
=====
Static Route Table (Router: Base)          Family : [IPv4|MCast-IPv4|IPv6]
=====
Network : 3FFD:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFE3/
120  Type : [Nexthop|Indirect|Blackhole]
Nexthop : [address | LSP label & name]      Nexthop type: [IP|LDP|RSVP-
TE]
```

```

Interface :
Metric : 1                      Prefence : 5
Active : [Y|N]                  Admin State : [Up|Down]
Tag :
BFD: [enable|disabled]

CPE-check: [enabled|disabled]    State: [Up|Down]
Target : <address>
Interval : [value | n/a]        Drop Count : <value>
Log : [Y|N]
CPE Host Up/Dn Time : 0d 16:32:28
CPE Echo Req Tx : 0             CPE Echo Reply Rx: 0
CPE Up Transitions : 0          CPE Down Transitions : 0
CPE TTL : 13
=====
A:sim1#

*A:CPM133>config>router# show router static-route 3.3.3.3/32 detail
=====
Static Route Table (Router: Base) Family: IPv4
=====
Prefix          : 3.3.3.3/32
Nexthop         : n/a
Type            : Blackhole
Interface       : n/a
Prefix List     : n/a
Metric          : 1
Admin State     : Up
BFD             : disabled
CPE-check       : disabled
Nexthop Type    : IP
Active          : Y
Prefix List Type : n/a
Preference      : 5
Tag             : 0
Community       : 100:33
-----
No. of Static Routes: 1
=====
*A:Dut-C> show router static-route 10.1.0.5/32 detail
=====
Static Route Table (Router: Base) Family: IPv4
=====
Prefix          : 10.1.0.5/32
Nexthop         : 1.0.0.2
Indirect        : Type
Interface       : n/a
Prefix List     : n/a
Metric          : 1
Source Class    : 0
Admin State     : Up
Creation Origin : manual
BFD             : disabled
Community       :
CPE-check       : disabled
Tunnel Resolution: filter
RSVP-TE Tunnels : enabled
Disallow-IGP    : disabled
LDP Tunnels     : disabled
-----
No. of Static Routes: 1
=====

```

service-prefix

| | |
|--------------------|---|
| Syntax | service-prefix |
| Context | show>router |
| Description | This command displays the address ranges reserved by this node for services sorted by prefix. |
| Output | Service Prefix Output — The following output is an example of service prefix information, and Table 31 describes the fields. |

Sample Output

```
A:ALA-A# show router service-prefix
=====
Address Ranges reserved for Services
=====
IP Prefix          Mask      Exclusive
-----
172.16.1.0         24       true
172.16.2.0         24       false
=====
A:ALA-A#
```

Table 31 **Service Prefix Fields**

| Label | Description |
|-----------|---|
| IP Prefix | The IP prefix of the range of addresses included in the range for services. |
| Mask | The subnet mask length associated with the IP prefix. |
| Exclusive | False — Addresses in the range are not exclusively for use for service IP addresses. True — Addresses in the range are exclusively for use for service IP addresses and cannot be assigned to network IP interfaces. |

sgt-qos

| | |
|--------------------|--|
| Syntax | sgt-qos |
| Context | show>router |
| Description | This command displays DSCP/dot1p remarking information for self-generated traffic. |

application

- Syntax** `application [app-name] [dscp | dot1p]`
- Context** `show>router>sgt-qos`
- Description** This command displays application QoS settings.
- Parameters** *app-name* — the specific application
- Values** arp, bgp, cflowd, dhcp, diameter, dns, ftp, gtp, icmp, igmp, igmp-reporter, isis, l2tp, ldp, mld, msdp, ndis, ntp, ospf, pcep, pim, pppoe, ptp, radius, rip, rsvp, sflow, snmp, snmp-notification, srrp, ssh, syslog, tacplus, telnet, tftp, traceroute, vrrp
- dscp** — Specifies to show DSCP values only.
- dot1p** — Specifies to show dot1p values only.
- Output** The following output is an example of SGT QoS application information.

Sample Output

```
A:ALA-A# show router sgt-qos application
=====
DSCP Application Values
=====
Application          Configured DSCP Value          Default DSCP Value(s)
-----
bgp                   none                            nc1
cflowd                none                            nc1
dhcp                  none                            nc1, af41, nc2
diameter              none                            af41
dns                   none                            af41
ftp                   none                            af41
gtp                   none                            nc2, nc1
icmp                  none                            be, nc1
igmp                  none                            nc1
igmp-reporter         none                            nc1
l2tp                  none                            nc1
ldp                   none                            nc1
mld                   none                            nc1
msdp                  none                            nc1
ndis                  none                            nc1, nc2
ntp                   none                            nc1
ospf                  none                            nc1
pcep                  none                            nc1
pim                   none                            nc1
ptp                   none                            nc1
radius                none                            nc1
rip                   none                            nc1
rsvp                  none                            nc1
sflow                 none                            nc1
snmp                  none                            af41
snmp-notification    none                            af41
srrp                  none                            nc1
ssh                   none                            af41
```

```

syslog          none          af41
tacplus         none          af41
telnet          none          af41
tftp            none          af41
traceroute      none          be
vrrp            none          nc1
=====
Dot1p Application Values
=====
Application      Configured Dot1p Value      Default Dot1p Value
-----
arp              none                          7
isis             none                          7
pppoe            none                          7
=====
A:ALA-A#

```



Note: Some applications have multiple DSCP default values depending on the context or service.

dscp-map

- Syntax** `dscp-map [dscp-name]`
- Context** `show>router>sgt-qos`
- Description** This command displays DSCP to FC mappings.
- Parameters** *dscp-name* — the specific DSCP name
- Values** be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63
- Output** The following output is an example of DSCP mapping information.

Sample Output

```

A:ALA-A# show router sgt-qos dscp-map
=====
DSCP to FC Mappings
=====
DSCP Value      FC Value      Default FC Value
-----
be              nc              nc
cp1             be              be
cp2             be              be

```

| | | |
|------|----|----|
| cp3 | be | be |
| cp4 | be | be |
| cp5 | be | be |
| cp6 | be | be |
| cp7 | be | be |
| cs1 | be | be |
| cp9 | be | be |
| af11 | af | af |
| cp11 | be | be |
| af12 | af | af |
| cp13 | be | be |
| af13 | af | af |
| cp15 | be | be |
| cs2 | be | be |
| cp17 | be | be |
| af21 | l1 | l1 |
| cp19 | be | be |
| af22 | l1 | l1 |
| cp21 | be | be |
| af23 | l1 | l1 |
| cp23 | be | be |
| cs3 | be | be |
| cp25 | be | be |
| af31 | l1 | l1 |
| cp27 | be | be |
| af32 | l1 | l1 |
| cp29 | be | be |
| af33 | l1 | l1 |
| cp31 | be | be |
| cs4 | be | be |
| cp33 | be | be |
| af41 | nc | nc |
| cp35 | be | be |
| af42 | h2 | h2 |
| cp37 | be | be |
| af43 | h2 | h2 |
| cp39 | be | be |
| cs5 | be | be |
| cp41 | be | be |
| cp42 | be | be |
| cp43 | be | be |
| cp44 | be | be |
| cp45 | be | be |
| ef | ef | ef |
| cp47 | be | be |
| nc1 | nc | nc |
| cp49 | be | be |
| cp50 | h2 | h2 |
| cp51 | be | be |
| cp52 | be | be |
| cp53 | be | be |
| cp54 | be | be |
| cp55 | be | be |
| nc2 | nc | nc |
| cp57 | be | be |
| cp58 | be | be |
| cp59 | be | be |
| cp60 | be | be |
| cp61 | be | be |


```

cp62                be                be
cp63                be                be
=====
A:ALA-A#

```

status

| | |
|--------------------|---|
| Syntax | status |
| Context | show>router |
| Description | This command displays the router status. |
| Output | Router Status Output — The following output is an example of router status information, and Table 32 describes the fields. |

There are multiple instances of OSPF. OSPF-0 is persistent. OSPF-1 through OSPF-31 are present when that specific OSPF instance is configured.

Sample Output

```

*A:Performance# show router status
=====
Router Status (Router: Base)
=====

```

| | Admin State | Oper State |
|------------------------|--------------------|----------------|
| Router | Up | Up |
| OSPFv2-0 | Up | Up |
| RIP | Up | Up |
| ISIS | Up | Up |
| MPLS | Not configured | Not configured |
| RSVP | Not configured | Not configured |
| LDP | Not configured | Not configured |
| BGP | Up | Up |
| IGMP | Not configured | Not configured |
| PIM | Not configured | Not configured |
| OSPFv3 | Not configured | Not configured |
| MSDP | Not configured | Not configured |
| Max Routes | No Limit | |
| Total IPv4 Routes | 244285 | |
| Total IPv6 Routes | 0 | |
| Max Multicast Routes | No Limit | |
| Total Multicast Routes | PIM not configured | |
| ECMP Max Routes | 1 | |
| Triggered Policies | No | |

```

=====
*A:Performance#

```

Table 32 Router Status Fields

| Label | Description |
|--------------------------|---|
| Router | The administrative and operational states for the router. |
| OSPF | The administrative and operational states for the OSPF protocol. |
| RIP | The administrative and operational states for RIP. |
| ISIS | The administrative and operational states for the IS-IS protocol. |
| MPLS | The administrative and operational states for the MPLS protocol. |
| RSVP | The administrative and operational states for RSVP. |
| LDP | The administrative and operational states for LDP. |
| BGP | The administrative and operational states for BGP. |
| IGMP | The administrative and operational states for IGMP. |
| MLD | The administrative and operational states for the MLD protocol. |
| PIM | The administrative and operational states for the PIM protocol. |
| OSPFv3 | The administrative and operational states for the OSPFv3 protocol. |
| MSDP | The administrative and operational states for MSDP. |
| Max Routes | The maximum number of routes configured for the system. |
| Total Routes | The total number of routes in the route table. |
| ECMP Max Routes | The number of ECMP routes configured for path sharing. |
| <i>service-id</i> | state — Current single SFM state start — Last time this vRtr went into overload, after having respected the hold-off time interval — How long the vRtr remained or is in overload |
| ICMP Tunneling | No — ICMP tunneling is disabled. Yes — TICMP tunneling is enabled. |
| VPRN Local TTL Propagate | inherit — VPRN instance is to inherit the global configuration none — TTL of IP packet is not propagated into the VC or transport label stack vc-only — TTL of the IP packet is propagated into the VC label and not into the labels in the transport label stack all — TTL of the IP packet is propagated into the VC label and all labels in the transport label stack |

Table 32 Router Status Fields (Continued)

| Label | Description |
|--------------------------|--|
| VPRN Transit TTL Propag* | inherit — VPRN instance is to inherit the global configuration none — TTL of IP packet is not propagated into the VC or transport label stack vc-only — TTL of the IP packet is propagated into the VC label and not into the labels in the transport label stack al — TTL of the IP packet is propagated into the VC label and all labels in the transport label stack |
| Label Route Local TTL P* | all — TTL of the IP packet is propagated into all labels of the transport label stack none — TTL of the IP packet is not propagated into the transport label stack |
| Label Route Transit TTL* | all — TTL of the IP packet is propagated into all labels of the transport label stack none — TTL of the IP packet is not propagated into the transport label stack |
| LSR Label Route TTL Pro* | all — TTL of the swapped label is propagated into all labels of the transport label stack none — TTL of the swapped label is not propagated into the transport label stack |
| Triggered Policies | No — Triggered route policy re-evaluation is disabled Yes — Triggered route policy re-evaluation is enabled |
| Class Forwarding | Enabled — Class Forwarding is enabled Disabled — Class Forwarding is disabled |

7450 ESS Router Status Output—The following output is an example of router status information for the 7450 ESS:

Sample Output

```
*A:Performance# configure router ospf [1..31] shutdown
*A:Performance# show router status
=====
Router Status (Router: Base)
=====
Admin State      Oper State
-----
Router           Up           Up
OSPFv2-0         Up           Up
OSPFv2-1         Down        Down
OSPFv2-2         Down        Down
```

```

OSPFv2-3           Down           Down
OSPFv2-4           Down           Down
OSPFv2-5           Down           Down
OSPFv2-6           Down           Down
OSPFv2-7           Down           Down
OSPFv2-8           Down           Down
OSPFv2-9           Down           Down
OSPFv2-10          Down           Down
OSPFv2-11          Down           Down
OSPFv2-12          Down           Down
OSPFv2-13          Down           Down
OSPFv2-14          Down           Down
OSPFv2-15          Down           Down
OSPFv2-16          Down           Down
OSPFv2-17          Down           Down
OSPFv2-18          Down           Down
OSPFv2-19          Down           Down
OSPFv2-20          Down           Down
OSPFv2-21          Down           Down
OSPFv2-22          Down           Down
OSPFv2-23          Down           Down
OSPFv2-24          Down           Down
OSPFv2-25          Down           Down
OSPFv2-26          Down           Down
OSPFv2-27          Down           Down
OSPFv2-28          Down           Down
OSPFv2-29          Down           Down
OSPFv2-30          Down           Down
OSPFv2-31          Down           Down
RIP                 Up             Up
ISIS                Up             Up
MPLS                Not configured Not configured
RSVP                 Not configured Not configured
LDP                  Not configured Not configured
BGP                  Up             Up
IGMP                 Not configured Not configured
PIM                  Not configured Not configured
OSPFv3               Not configured Not configured
MSDP                 Not configured Not configured
Max Routes           No Limit
Total IPv4 Routes    244277
Max Multicast Routes No Limit
Total Multicast Routes PIM not configured
ECMP Max Routes      1
Single SFM Overload   Enabled           hold-off 30 sec
Single SFM State      normal
Single SFM Start      004 19:03:39.680
Single SFM Interval   0d 00:16:06
Reassembly ISA-BB group Not configured
Ipv6 Nbr Reachab. time Not configured      30
Triggered Policies    No
=====
*A:Performance#

```

Router Status Output for 7750 SR and 7950 XRS—The following output is an example of router status information for the 7750 SR and 7950 XRS:

Sample Output

```
*A:Performance# configure router ospf [1..31] shutdown
*A:Performance# show router status
=====
Router Status (Router: Base)
=====
-----
Admin State      Oper State
-----
Router           Up            Up
OSPFv2-0         Up            Up
OSPFv2-1         Down          Down
OSPFv2-2         Down          Down
OSPFv2-3         Down          Down
OSPFv2-4         Down          Down
OSPFv2-5         Down          Down
OSPFv2-6         Down          Down
OSPFv2-7         Down          Down
OSPFv2-8         Down          Down
OSPFv2-9         Down          Down
OSPFv2-10        Down          Down
OSPFv2-11        Down          Down
OSPFv2-12        Down          Down
OSPFv2-13        Down          Down
OSPFv2-14        Down          Down
OSPFv2-15        Down          Down
OSPFv2-16        Down          Down
OSPFv2-17        Down          Down
OSPFv2-18        Down          Down
OSPFv2-19        Down          Down
OSPFv2-20        Down          Down
OSPFv2-21        Down          Down
OSPFv2-22        Down          Down
OSPFv2-23        Down          Down
OSPFv2-24        Down          Down
OSPFv2-25        Down          Down
OSPFv2-26        Down          Down
OSPFv2-27        Down          Down
OSPFv2-28        Down          Down
OSPFv2-29        Down          Down
OSPFv2-30        Down          Down
OSPFv2-31        Down          Down
RIP              Up            Up
ISIS             Up            Up
MPLS             Not configured Not configured
RSVP             Not configured Not configured
LDP              Not configured Not configured
BGP              Up            Up
IGMP             Not configured Not configured
PIM              Not configured Not configured
OSPFv3           Not configured Not configured
MSDP             Not configured Not configured
Max Routes       No Limit
Total IPv4 Routes 244277
Total IPv6 Routes 0
Max Multicast Routes No Limit
Total Multicast Routes PIM not configured
ECMP Max Routes 1
Single SFM Overload Enabled          hold-off 30 sec
```

```

Single SFM State          normal
Single SFM Start          004 19:03:39.680
Single SFM Interval       0d 00:16:06
Reassembly ISA-BB group   Not configured
Ipv6 Nbr Reachab. time    Not configured          30
Triggered Policies        No
=====
*A:Performance#

```

Class Forwarding—The following output is an example for checking if class-based forwarding is enabled in the global router context.

Sample Output

```

*A:Dut-B>show>router# show router "Base" status
=====
Router Status (Router: Base)
=====

```

| | Admin State | Oper State |
|-------------------------|--------------------|----------------|
| Router | Up | Up |
| OSPFv2 | Not configured | Not configured |
| RIP | Not configured | Not configured |
| RIP-NG | Not configured | Not configured |
| ISIS-0 | Up | Up |
| MPLS | Up | Up |
| RSVP | Up | Up |
| LDP | Not configured | Not configured |
| BGP | Up | Up |
| IGMP | Not configured | Not configured |
| MLD | Not configured | Not configured |
| PIM | Not configured | Not configured |
| PIMv4 | Not configured | Not configured |
| PIMv6 | Not configured | Not configured |
| OSPFv3 | Not configured | Not configured |
| MSDP | Not configured | Not configured |
| Max IPv4 Routes | No Limit | |
| Max IPv6 Routes | No Limit | |
| Total IPv4 Routes | 262 | |
| Total IPv6 Routes | 262 | |
| Max Multicast Routes | No Limit | |
| Total IPv4 Mcast Routes | PIM not configured | |
| Total IPv6 Mcast Routes | PIM not configured | |
| ECMP Max Routes | 64 | |
| Weighted ECMP | Disabled | |
| Mcast Info Policy | default | |
| Triggered Policies | No | |
| LDP Shortcut | Disabled | |
| Single SFM Overload | Disabled | |
| IP Fast Reroute | Disabled | |
| Entropy Label | Disabled | |
| ICMP Tunneling | Enabled | |
| MSS adjust ISA group | Not configured | |
| Reassembly ISA-BB group | Not configured | |
| Ipv6 Nbr Reachab. time | Not configured | 30 |
| IPv6 Nbr stale time (s) | 14400 | |
| Static Route Hold Down | Disabled | |

```
TTL Propagate
  VPRN Local      vc-only
  VPRN Transit    vc-only
  Label Route Local none
  Label Route Transit none
  LSR Label Route none
LSP BFD Tail Sessions Disabled
Class Forwarding  Enabled
=====
```

TTL Propagation and ICMP Tunneling—The following output is an example of TTL propagation and ICMP tunneling configurations, first in base router and then in a VPRN service.

Sample Output

```
*A:Performance# show router status
=====
Router Status (Router: Base)
=====
```

| | Admin State | Oper State |
|-----------|----------------|----------------|
| Router | Up | Up |
| OSPFv2-0 | Up | Up |
| OSPFv2-2 | Down | Down |
| RIP | Not configured | Not configured |
| RIP-NG | Not configured | Not configured |
| ISIS-0 | Up | Up |
| ISIS-1024 | Down | Down |
| MPLS | Down | Down |
| RSVP | Down | Down |
| LDP | Up | Down |
| BGP | Up | Down |
| IGMP | | |
| MLD | | |
| PIM | | |
| PIMv4 | | |
| PIMv6 | | |
| OSPFv3 | | |
| MSDP | | |

```

Max IPv4 Routes      No Limit
Max IPv6 Routes      No Limit
Total IPv4 Routes    0
Total IPv6 Routes    0
Max Multicast Routes No Limit
Total IPv4 Mcast Routes PIM not configured
Total IPv6 Mcast Routes PIM not configured
ECMP Max Routes      1
Mcast Info Policy    default
Triggered Policies   No
LDP Shortcut         Disabled
Single SFM Overload  Disabled
IP Fast Reroute      Disabled
ICMP Tunneling       Disabled
Reassembly ISA-BB group Not configured
ICMP Tunneling       Disabled

```

```

Ipv6 Nbr Reachab. time    Not configured          30
IPv6 Nbr stale time (s)   14400
VPRN Local TTL Propagate  vc-only
VPRN Transit TTL Propag* vc-only
Label Route Local TTL P* none
Label Route Transit TTL* none
LSR Label Route TTL Pro* none
=====
* indicates that the corresponding row element may have been truncated.
*B:bkvm31#

```

VPRN TTL Propagation and ICMP Tunneling—The following output is an example of TTL propagation and ICMP tunneling configurations in a VPRN service. The ttl-propagation has been specified as local and all for VPRN service 5001.

Sample Output

```

*A:Dut-A# configure service vprn 5001 ttl-propagate local all
*A:Dut-A# show router 5001 status

```

```

=====
Router Status (Service: 5001)
=====

```

| | Admin State | Oper State |
|-------------------------|--------------------|----------------|
| Router | Up | Up |
| OSPFv2 | Not configured | Not configured |
| RIP | Not configured | Not configured |
| RIP-NG | Not configured | Not configured |
| ISIS | Not configured | Not configured |
| MPLS | Not configured | Not configured |
| RSVP | Not configured | Not configured |
| LDP | Not configured | Not configured |
| BGP | Not configured | Not configured |
| IGMP | Not configured | Not configured |
| MLD | Not configured | Not configured |
| PIM | Not configured | Not configured |
| PIMv4 | Not configured | Not configured |
| PIMv6 | Not configured | Not configured |
| OSPFv3 | Not configured | Not configured |
| MSDP | Not configured | Not configured |
| Max IPv4 Routes | No Limit | |
| Max IPv6 Routes | No Limit | |
| Total IPv4 Routes | 2 | |
| Total IPv6 Routes | 2 | |
| Max Multicast Routes | No Limit | |
| Total IPv4 Mcast Routes | PIM not configured | |
| Total IPv6 Mcast Routes | PIM not configured | |
| ECMP Max Routes | 1 | |
| Mcast Info Policy | default | |
| Triggered Policies | No | |
| GRT Lookup | Disabled | |
| Local Management | Disabled | |
| Single SFM Overload | Disabled | |
| IP Fast Reroute | Disabled | |


```

ICMP Tunneling           Disabled
Reassembly ISA-BB group  Not configured
ICMP Tunneling           Disabled
Ipv6 Nbr Reachab. time   Not configured          30
VPRN Local TTL Propagate all
VPRN Transit TTL Propag* inherit (vc-only)
=====
* indicates that the corresponding row element may have been truncated.
*A:Dut-A#

```

tunnel-table

```
Syntax  tunnel-table summary [ipv4 | ipv6]
        tunnel-table [protocol protocol] {ipv4 | ipv6}
        tunnel-table [ip-prefix[/mask]] [alternative] [ipv4 | ipv6] [detail]
        tunnel-table mpls-tp
        tunnel-table [ip-prefix[/mask]] protocol protocol [detail]
        tunnel-table [ip-prefix[/mask]] sdp sdp-id
```

Context show>router

| | |
|--------------------|---|
| Description | This command displays tunnel table information. Auto-bind GRE tunnels are not displayed in show command output. GRE tunnels are not the same as SDP tunnels that use the GRE encapsulation type. When the auto-bind command is used when configuring a VPRN service, it means the MP-BGP NH resolution is referring to the core routing instance for IP reachability. For a VPRN service this object specifies the lookup to be used by the routing instance if no SDP to the destination exists. |
|--------------------|---|

Parameters *ip-address[/mask]* — displays the specified tunnel table's destination IP address and mask

protocol *protocol* — displays protocol information

| | |
|---------------|---|
| Values | bgp, ldp, rsvp, sdp, ospf, isis, sr-te, fpe |
|---------------|---|

sdp *sdp-id* — displays information pertaining to the specified SDP

Values 1 to 17407

summary — displays summary tunnel table information

detail — displays detailed tunnel table information

alternative — displays Backup Route details

mpls-tp — displays MPLS-TP information

ipv4 — displays information for IPv4 entries only

ipv6 — displays information for IPv6 entries only

Output **Tunnel Table Output** — The following output is an example of tunnel table information, and [Table 33](#) describes the fields.

Sample Output

```

*A:Dut-D>config>service>vpls# show router tunnel-table sdp 17407
=====
Tunnel Table (Router: Base)
=====
Destination          Owner Encap TunnelId Pref    Nexthop      Metric
-----
127.0.68.0/32        sdp   MPLS   17407    5        127.0.68.0    0
=====
*A:Dut-D# show service id 1 sdp 17407:4294967294 detail
=====
Service Destination Point (Sdp Id : 17407:4294967294) Details
=====
-----
Sdp Id 17407:4294967294  -(not applicable)
-----
Description           : (Not Specified)
SDP Id                 : 17407:4294967294          Type                : VplsPmsi
Split Horiz Grp        : (Not Specified)
VC Type                : Ether                      VC Tag              : n/a
Admin Path MTU         : 9194                      Oper Path MTU       : 9194
Delivery               : MPLS
Far End                : not applicable
Tunnel Far End         : n/a                      LSP Types           : None
Hash Label             : Disabled                  Hash Lbl Sig Cap    : Disabled
Oper Hash Label        : Disabled

Admin State            : Up                      Oper State          : Up
Acct. Pol              : None                    Collect Stats       : Disabled
Ingress Label          : 0                      Egress Label       : 3
Ingr Mac Fltr-Id       : n/a                    Egr Mac Fltr-Id    : n/a
Ingr IP Fltr-Id        : n/a                    Egr IP Fltr-Id     : n/a
Ingr IPv6 Fltr-Id      : n/a                    Egr IPv6 Fltr-Id   : n/a
Admin ControlWord      : Not Preferred    Oper ControlWord    : False
Last Status Change     : 12/14/2012 12:42:22        Signaling           : None
Last Mgmt Change       : 12/14/2012 12:42:19        Force Vlan-Vc       : Disabled
Endpoint               : N/A                      Precedence          : 4
PW Status Sig          : Enabled
Class Fwding State     : Down
Flags                  : None
Time to RetryReset     : never                      Retries Left        : 3
Mac Move               : Blockable                  Blockable Level     : Tertiary
Local Pw Bits          : None
Peer Pw Bits           : None
Peer Fault Ip          : None
Peer Vccv CV Bits      : None
Peer Vccv CC Bits      : None
Application Profile    : None
Max Nbr of MAC Addr    : No Limit                  Total MAC Addr      : 0
Learned MAC Addr       : 0                      Static MAC Addr     : 0

MAC Learning           : Enabled                  Discard Unkwn Srce  : Disabled
MAC Aging              : Enabled
BPDU Translation       : Disabled
L2PT Termination       : Disabled
MAC Pinning            : Disabled
Ignore Standby Sig     : False                  Block On Mesh Fail  : False
Oper Group              : (none)                  Monitor Oper Grp    : (none)

```

```

Rest Prot Src Mac : Disabled
Auto Learn Mac Prot: Disabled
RestProtSrcMacAct : Disable

Ingress Qos Policy : (none)
Ingress FP QGrp : (none)
Ing FP QGrp Inst : (none)
Egress Qos Policy : (none)
Egress Port QGrp : (none)
Egr Port QGrp Inst: (none)

-----
ETH-CFM SDP-Bind specifics
-----
V-MEP Filtering : Disabled

KeepAlive Information :
Admin State : Disabled
Hello Time : 10
Max Drop Count : 3
Oper State : Disabled
Hello Msg Len : 0
Hold Down Time : 10

Statistics :
I. Fwd. Pkts. : 0
I. Fwd. Octs. : 0
E. Fwd. Pkts. : 2979761
I. Dro. Pkts. : 0
I. Dro. Octs. : 0
E. Fwd. Octets : 476761760
-----
Control Channel Status
-----
PW Status : disabled
Peer Status Expire : false
Refresh Timer : <none>
Clear On Timeout : true

MCAC Policy Name :
MCAC Max Unconst BW: no limit
MCAC In use Mand BW: 0
MCAC In use Opnl BW: 0
MCAC Max Mand BW : no limit
MCAC Avail Mand BW: unlimited
MCAC Avail Opnl BW: unlimited
-----
RSVP/Static LSPs
-----
Associated LSP List :
No LSPs Associated
-----
Class-based forwarding :
-----
Class forwarding : Disabled
Default LSP : Uknwn
EnforceDSTELspFc : Disabled
Multicast LSP : None
=====
FC Mapping Table
=====
FC Name LSP Name
-----
No FC Mappings
-----
Stp Service Destination Point specifics
-----
Stp Admin State : Down
Core Connectivity : Down
Port Role : N/A
Port Number : 0
Port Path Cost : 10
Admin Edge : Disabled
Link Type : Pt-pt
Root Guard : Disabled
Last BPDU from : N/A
Stp Oper State : Down
Port State : Forwarding
Port Priority : 128
Auto Edge : Enabled
Oper Edge : N/A
BPDU Encap : Dot1d
Active Protocol : N/A

```

```
Designated Bridge : N/A                      Designated Port Id: N/A

Fwd Transitions   : 0                      Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd    : 0                      Cfg BPDUs tx     : 0
TCN BPDUs rcvd    : 0                      TCN BPDUs tx     : 0
TC bit BPDUs rcvd : 0                      TC bit BPDUs tx  : 0
RST BPDUs rcvd    : 0                      RST BPDUs tx     : 0
```

```
-----
Number of SDPs : 1
-----
=====
```

```
*A:Dut-C# show router tunnel-table sdp 17407
```

```
=====
Tunnel Table (Router: Base)
```

```
=====
Destination      Owner Encap TunnelId Pref  Nexthop      Metric
-----
127.0.68.0/32    sdp  MPLS  17407   5    127.0.68.0    0
=====
```

```
A:ALA-A>config>service# show router tunnel-table Tunnel Table
```

```
=====
DestinationOwnerEncapTunnel IdPrefNexthopMetric
-----
```

```
10.0.0.1/32 sdp GRE 10 5 10.0.0.1 0
10.0.0.1/32 sdp GRE 21 5 10.0.0.1 0
10.0.0.1/32 sdp GRE 31 5 10.0.0.1 0
10.0.0.1/32 sdp GRE 41 5 10.0.0.1 0
=====
```

```
A:ALA-A>config>service#
```

```
A:ALA-A>config>service# show router tunnel-table summary
```

```
=====
Tunnel Table Summary (Router: Base)
```

```
=====
Active Available
-----
LDP          1          1
SDP          1          1
=====
```

```
A:ALA-A>config>service#
```

```
A:Dut-C# show router tunnel-table
```

```
=====
Tunnel Table (Router: Base)
```

```
=====
Destination      Owner      Encap TunnelId Pref  Nexthop      Metric
-----
4.0.0.1/32       isis (0)   MPLS  524309   11    1.3.4.4      10
10.20.1.2/32     isis (0)   MPLS  524312   11    1.2.3.2      10
10.20.1.4/32     isis (0)   MPLS  524310   11    1.3.4.4      10
10.20.1.5/32     isis (0)   MPLS  524311   11    1.2.3.2      20
=====
```

```
Flags: B = BGP backup route available
```

```
E = inactive best-
external BGP route
=====
```

```
A:Dut-C#
```

```
*A:Dut-C> show router tunnel-table
```

```
=====
IPv4 Tunnel Table (Router: Base)
=====
```

| Destination | Owner | Encap | TunnelId | Pref | NextHop | Metric |
|--------------|----------|-------|----------|------|------------|--------|
| 10.20.1.1/32 | ldp | MPLS | 65546 | 9 | 10.10.2.1 | 10 |
| 10.20.1.2/32 | ldp | MPLS | 65545 | 9 | 10.10.12.2 | 3 |
| 10.20.1.2/32 | isis (0) | MPLS | 524318 | 11 | 10.10.12.2 | 3 |
| 10.20.1.4/32 | isis (0) | MPLS | 524316 | 11 | 10.10.11.4 | 10 |
| 10.20.1.5/32 | ldp | MPLS | 65547 | 9 | 10.10.5.5 | 10 |
| 10.20.1.5/32 | isis (0) | MPLS | 524315 | 11 | 10.10.5.5 | 10 |
| 10.20.1.6/32 | isis (0) | MPLS | 524317 | 11 | 10.10.11.4 | 20 |

```
-----
Flags: B = BGP backup route available
      E = inactive best-external BGP route
=====
```

```
A:Dut-C# show router tunnel-table detail
```

```
=====
Tunnel Table (Router: Base)
=====
```

```
Destination      : 7.1.126.2/32
NextHop          : 110.20.1.5
Tunnel Flags     : is-over-tunnel
Age              : 01h27m59s
CBF Classes      : (Not Specified)
Owner            : ldp
Tunnel ID        : 66389
Tunnel Label     : 243909
Tunnel MTU       : 9186
Encap            : MPLS
Preference       : 9
Tunnel Metric    : 1
-----
```

```
Destination      : 10.20.1.22/32
NextHop          : 120.1.17.7
Tunnel Flags     : (Not Specified)
Age              : 01h29m15s
CBF Classes      : (Not Specified)
Owner            : rsvp
Tunnel ID        : 13
Tunnel Label     : 249809
Tunnel MTU       : 9190
LSP ID           : 44032
LSP Bandwidth    : 0
Encap            : MPLS
Preference       : 7
Tunnel Metric    : 2000
Bypass Label     : 0
LSP Weight       : 0
-----
```

```
Destination      : 10.20.1.22/32
NextHop          : 120.1.18.7
Tunnel Flags     : exclude-for-lfa
Age              : 00h01m47s
CBF Classes      : default-lsp
Owner            : rsvp
Tunnel ID        : 243
Tunnel Label     : 249872
Tunnel MTU       : 9190
LSP ID           : 44032
LSP Bandwidth    : 0
Encap            : MPLS
Preference       : 7
Tunnel Metric    : 2000
Bypass Label     : 0
LSP Weight       : 0
-----
```

```

Destination      : 10.20.1.22/32
NextHop          : 120.1.18.7
Tunnel Flags     : exclude-for-lfa
Age              : 00h00m38s
CBF Classes      : af ll ef nc
Owner            : rsvp
Tunnel ID        : 244
Tunnel Label     : 249905
Tunnel MTU       : 9190
LSP ID           : 45568
LSP Bandwidth    : 0
Encap            : MPLS
Preference       : 7
Tunnel Metric    : 2000
Bypass Label     : 0
LSP Weight       : 0

```

```

-----
Destination      : 10.20.1.22/32
NextHop          : 120.1.17.7
Tunnel Flags     : exclude-for-lfa
Age              : 00h00m21s
CBF Classes      : h2
Owner            : rsvp
Tunnel ID        : 245
Tunnel Label     : 250063
Tunnel MTU       : 9190
LSP ID           : 39936
LSP Bandwidth    : 0
Encap            : MPLS
Preference       : 7
Tunnel Metric    : 2000
Bypass Label     : 0
LSP Weight       : 0

```

```

-----
Destination      : 10.20.1.22/32
NextHop          : 120.1.18.7
Tunnel Flags     : exclude-for-lfa
Age              : 01h29m40s
CBF Classes      : ef default-lsp
Owner            : rsvp
Tunnel ID        : 246
Tunnel Label     : 250024
Tunnel MTU       : 9190
LSP ID           : 38400
LSP Bandwidth    : 0
Encap            : MPLS
Preference       : 7
Tunnel Metric    : 2000
Bypass Label     : 0
LSP Weight       : 0

```

```

-----
Destination      : 211.1.0.254/32
NextHop          : 110.20.1.4
Tunnel Flags     : is-over-tunnel
Age              : 01h28m38s
CBF Classes      : (Not Specified)
Owner            : bgp
Tunnel ID        : 264115
Tunnel Label     : 260512
Tunnel MTU       : 9186
Encap            : MPLS
Preference       : 12
Tunnel Metric    : 1000

```

```

-----
Number of tunnel-table entries      : 2866
Number of tunnel-table entries with LFA : 0
=====

```

A:Dut-C#

*B:Dut-B>config>router>mpls>lsp# show router tunnel-table ipv6 protocol isis

=====

IPv6 Tunnel Table (Router: Base)

```

=====
Destination      Owner      Encap  TunnelId  Pref
NextHop          Metric
=====

```

```

2001::a14:103/128          isis (0)  MPLS  524355    11
fe80::c7b:1ff:fe01:1-"B_to_D"          20
2001::a14:104/128          isis (0)  MPLS  524354    11
fe80::c7b:1ff:fe01:1-"B_to_D"          10
2001::a14:105/128          isis (0)  MPLS  524356    11
fe80::c7f:2ff:fe01:1-"B_to_E"          10
2001::a14:106/128          isis (0)  MPLS  524357    11
fe80::c7b:1ff:fe01:1-"B_to_D"          20

```

Flags: B = BGP backup route available
E = inactive best-external BGP route
=====

*B:Dut-B>config>router>mpls>lsp# show router tunnel-table ipv6 detail

=====

Tunnel Table (Router: Base)

=====

```

Destination      : 2001::a14:103/128
NextHop          : 2001::a14:103
Tunnel Flags     : (Not Specified)
Age              : 00h02m20s
CBF Classes     : (Not Specified)
Owner            : sdp
Encap            : MPLS
Tunnel ID        : 230
Preference       : 5
Tunnel Label     : -
Tunnel Metric    : 0
Tunnel MTU       : 1578
Max Label Stack  : 1

```

```

Destination      : 2001::a14:103/128
NextHop          : fe80::c7b:1ff:fe01:1-"B_to_D"
Tunnel Flags     : (Not Specified)
Age              : 00h02m15s
CBF Classes     : (Not Specified)
Owner            : ldp
Encap            : MPLS
Tunnel ID        : 65567
Preference       : 9
Tunnel Label     : 262136
Tunnel Metric    : 200
Tunnel MTU       : 1582
Max Label Stack  : 1

```

```

Destination      : 2001::a14:103/128
NextHop          : fe80::c7b:1ff:fe01:1-"B_to_D"
Tunnel Flags     : exclude-for-igpshortcuts
Age              : 00h02m23s
CBF Classes     : (Not Specified)
Owner            : isis (0)
Encap            : MPLS
Tunnel ID        : 524355
Preference       : 11
Tunnel Label     : 18563
Tunnel Metric    : 20
Tunnel MTU       : 1582
Max Label Stack  : 1

```

```

Destination      : 2001::a14:104/128
NextHop          : fe80::c7b:1ff:fe01:1-"B_to_D"
Tunnel Flags     : (Not Specified)
Age              : 00h02m20s
CBF Classes     : (Not Specified)
Owner            : ldp
Encap            : MPLS
Tunnel ID        : 65568
Preference       : 9
Tunnel Label     : 262143
Tunnel Metric    : 100
Tunnel MTU       : 1582
Max Label Stack  : 1

```

```

Destination      : 2001::a14:104/128
NextHop          : fe80::c7b:1ff:fe01:1-"B_to_D"

```

```

Tunnel Flags      : exclude-for-igpshortcuts
Age               : 00h02m32s
CBF Classes       : (Not Specified)
Owner            : isis (0)           Encap             : MPLS
Tunnel ID         : 524354           Preference        : 11
Tunnel Label      : 18564           Tunnel Metric     : 10
Tunnel MTU        : 1582           Max Label Stack  : 1
-----
Destination       : 2001::a14:105/128
NextHop           : fe80::c7f:2ff:fe01:1-"B_to_E"
Tunnel Flags      : (Not Specified)
Age               : 00h02m15s
CBF Classes       : (Not Specified)
Owner            : ldp               Encap             : MPLS
Tunnel ID         : 65569           Preference        : 9
Tunnel Label      : 262143          Tunnel Metric     : 100
Tunnel MTU        : 1582           Max Label Stack  : 1
-----
Destination       : 2001::a14:105/128
NextHop           : fe80::c7f:2ff:fe01:1-"B_to_E"
Tunnel Flags      : exclude-for-igpshortcuts
Age               : 00h02m32s
CBF Classes       : (Not Specified)
Owner            : isis (0)           Encap             : MPLS
Tunnel ID         : 524356           Preference        : 11
Tunnel Label      : 18565           Tunnel Metric     : 10
Tunnel MTU        : 1582           Max Label Stack  : 1
-----
Destination       : 2001::a14:106/128
NextHop           : fe80::c7b:1ff:fe01:1-"B_to_D"
Tunnel Flags      : (Not Specified)
Age               : 00h02m16s
CBF Classes       : (Not Specified)
Owner            : ldp               Encap             : MPLS
Tunnel ID         : 65570           Preference        : 9
Tunnel Label      : 262133          Tunnel Metric     : 200
Tunnel MTU        : 1582           Max Label Stack  : 1
-----
Destination       : 2001::a14:106/128
NextHop           : fe80::c7b:1ff:fe01:1-"B_to_D"
Tunnel Flags      : exclude-for-igpshortcuts
Age               : 00h02m24s
CBF Classes       : (Not Specified)
Owner            : isis (0)           Encap             : MPLS
Tunnel ID         : 524357           Preference        : 11
Tunnel Label      : 18566           Tunnel Metric     : 20
Tunnel MTU        : 1582           Max Label Stack  : 1
-----
Number of tunnel-table entries      : 9
Number of tunnel-table entries with LFA : 0
=====

```

Table 33 Tunnel Table Fields

| Label | Description |
|-------------|---|
| Destination | The route's destination address and mask. |

Table 33 Tunnel Table Fields (Continued)

| Label | Description |
|-------------|--|
| Owner | Specifies the tunnel owner. |
| Encap | Specifies the tunnel's encapsulation type. |
| Tunnel ID | Specifies the tunnel (SDP) identifier. |
| Pref | Specifies the route preference for routes learned from the configured peers. |
| Nexthop | The next hop for the route's destination. |
| Metric | The route metric value for the route. |
| CBF Classes | The forwarding classes and/or default-lsp option assigned to this tunnel. |

CBF Info—The following output is an example for checking the CBF info of a tunnel in TTM.

Sample Output

```
*A:Dut-B>show>router# show router tunnel-table 10.20.1.5/32 protocol rsvp detail
=====
Tunnel Table (Router: Base)
=====
Destination      : 10.20.1.5/32
NextHop          : 10.11.18.3
Tunnel Flags     : exclude-for-lfa entropy-label-capable
Age              : 00h17m17s
CBF Classes      : (Not Specified)
Owner            : rsvp                               Encap            : MPLS
Tunnel ID        : 35                               Preference       : 7
Tunnel Label     : 262142                           Tunnel Metric    : 2000
Tunnel MTU       : 1496                             Max Label Stack  : 1
LSP ID           : 15872                             Bypass Label     : 0
LSP Bandwidth    : 0                               LSP Weight       : 0
-----
.
.
-----
Destination      : 10.20.1.5/32
NextHop          : 10.11.20.3
Tunnel Flags     : exclude-for-lfa entropy-label-capable
Age              : 00h17m18s
CBF Classes      : be l2
Owner            : rsvp                               Encap            : MPLS
Tunnel ID        : 47                               Preference       : 7
Tunnel Label     : 262070                           Tunnel Metric    : 2000
Tunnel MTU       : 1496                             Max Label Stack  : 1
LSP ID           : 15360                             Bypass Label     : 0
LSP Bandwidth    : 0                               LSP Weight       : 0
-----
.
```

```

.
.
-----
Destination      : 10.20.1.5/32
NextHop          : 10.11.7.3
Tunnel Flags     : exclude-for-lfa entropy-label-capable
Age              : 00h17m18s
CBF Classes      : af ll
Owner            : rsvp
Tunnel ID        : 60
Tunnel Label     : 262063
Tunnel MTU       : 1496
LSP ID           : 16384
LSP Bandwidth    : 0
Encap            : MPLS
Preference       : 7
Tunnel Metric    : 2000
Max Label Stack  : 1
Bypass Label     : 0
LSP Weight       : 0
-----
.
.
.
-----
Destination      : 10.20.1.5/32
NextHop          : 10.11.1.3
Tunnel Flags     : exclude-for-lfa entropy-label-capable
Age              : 00h17m19s
CBF Classes      : h2 ef default-lsp
Owner            : rsvp
Tunnel ID        : 76
Tunnel Label     : 262053
Tunnel MTU       : 1496
LSP ID           : 36352
LSP Bandwidth    : 0
Encap            : MPLS
Preference       : 7
Tunnel Metric    : 2000
Max Label Stack  : 1
Bypass Label     : 0
LSP Weight       : 0
-----
.
.
.
-----
Number of tunnel-table entries      : 64
Number of tunnel-table entries with LFA : 0
=====

```

2.13.2.1.1 L2TP Show Commands

The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

l2tp

| | |
|--------------------|---|
| Syntax | l2tp |
| Context | show>router |
| Description | This command enables the context to display L2TP related information. |

eth-tunnel

| | |
|--------------------|--|
| Syntax | eth-tunnel [group <i>tunnel-group-name</i> [vc-id <i>vc-id</i>]] |
| Context | show>router>l2tp |
| Description | <p>This command displays information about configured L2TPv3 Ethernet tunnels. These Ethernet tunnels are the L2TPv3 sessions setup between the local private L2 SAP and the far end device.</p> <p>If this command is executed without any parameters, then a list of all configured Ethernet tunnels are displayed.</p> <p>If this command is executed with a tunnel group name or a VC-ID, then a detailed view of the associated Ethernet tunnel is displayed.</p> |
| Parameters | <p><i>tunnel-group-name</i> — specifies the configured tunnel group name used for the associated Ethernet tunnel</p> <p><i>vc-id</i> — specifies the VC-ID for the L2TPv3 Ethernet tunnel</p> <p>Values 0 to 4294967295</p> |
| Output | The following output is an example of L2TPv3 Ethernet tunnel information |

Sample Output

```

A:Dut-A# show router 200 l2tp eth-tunnel
=====
L2TPv3 Ethernet Tunnel Summary
=====
Tunnel Group name                               VC ID
-----
v3-group-1                                     100
-----
No. of ethernet tunnels: 1
=====

A:Dut-A# show router 200 l2tp eth-tunnel group "v3-group-1"
=====
L2TPv3 Ethernet Tunnel Status
=====
Group Name           : v3-group-1
VC ID                : 100
Local Conn ID        : 221122308
Ctrl Conn ID         : 221118464
Matches Cfg          : true
Down Reason          : N/A
Reconnect Time (s)   : N/A
Remaining Time (s)   : N/A
SAP ID               : tunnel-1.private:100
SAP Service ID       : 100
-----
No. of ethernet tunnels: 1
=====

```

group

- Syntax** **group** [*tunnel-group-name* [**statistics**]]
- Context** show>router>l2tp
- Description** This command displays L2TP group operational information.
- Parameters** *tunnel-group-name* — displays information for the specified tunnel group
statistics — displays statistics for the specified tunnel group
- Output** The following output is an example of L2TP group operational information.

Sample Output

```
*A:Dut-C# show router l2tp group
=====
L2TP Groups
=====
Group Name          Ses Limit Ses Assign      State  Tun Active Ses Active
                                   Tun Total  Ses Total
-----
isp1.group-1
                131071    existingFirst active      1          1
                                   1          1
isp1.group-2
                131071    weighted      active      2          5
                                   3          8
-----
No. of L2TP Groups: 2
=====
*A:Dut-C#

*A:Dut-C# show router l2tp group isp1.group-2
=====
Group Name: isp1.group-2
=====
Conn ID              Loc-Tu-ID Rem-Tu-ID State              Ses Active
  Group                               Ses Total
  Assignment
-----
143523840            2190      17525    established        2
  isp1.group-2                               3
    isp1.tunnel-3
236912640            3615      58919    closedByPeer        0
  isp1.group-2                               2
    isp1.tunnel-2
658178048            10043     33762    draining            3
  isp1.group-2                               3
    isp1.tunnel-2
-----
No. of tunnels: 3
=====
*A:Dut-C#

*A:Dut-C# show router l2tp group isp1.group-2 statistics
```

```

Group Name: ispl.group-2
-----
              Attempts    Failed    Failed-Aut              Active    Total
-----
Tunnels       3           0         0              2         3
Sessions      8           0        N/A              5         8
-----
              Pkt-Ctl              Pkt-Err              Octets
-----
Rx             51                  0              1224
Tx             51                  0              2796
-----
*A:Dut-C#

```

peer

Syntax **peer** *ip-address* [**statistics**] [{**udp-port** *port* | **ip**}]
peer [**draining**] [{**blacklisted** | **selectable** | **unreachable**}]

Context show>router>l2tp

Description This command displays information regarding all configured L2TP peers.

If this command is executed without specifying a peer IP address, then a list of all L2TP peers are listed along with the type of transport used and statistics on the total number of tunnels and sessions, as well as the number of active tunnels and sessions.

If this command is executed with a specific peer IP address, than a detailed view for that peer is displayed.

Parameters *ip-address* — specifies the L2TP peer address
statistics — displays the statistics for the given IP address
port — specifies the UDP port for the L2TP peer. This parameter is only supported with L2TPv2 peers.
ip — displays peers using IP transport
draining — displays only peers with draining tunnels
blacklisted — displays peers that are blacklisted
selectable — displays peers that are selectable
unreachable — displays peers that are deemed unreachable

Output The following output is an example of L2TP peer operational information.

Sample Output

```

A:Dut-A# show router 200 l2tp peer
=====
L2TP Peers

```

```

=====
Peer IP                               Port  Tun Active Ses Active
                               Drain Reachability Tun Total  Ses Total
-----
10.1.1.2                             ip    1      1      1
                               1      1      1
-----
No. of peers: 1
=====

```

A:Dut-A# show router 200 l2tp peer 10.1.1.2 ip

```

=====
Peer IP: 10.1.1.2
=====
Roles capab/actual: LAC LNS / - -      Draining      : false
Tunnels           : 1                  Tunnels Active : 1
Sessions          : 1                  Sessions Active: 1
Reachability      : reachable          Time Unreachable: N/A
=====
Conn ID   Loc-Tu-ID Rem-Tu-ID State          Blacklist-state  Ses Active
  Group                                     Ses Total
  Assignment
-----
221118464 3374      0      established      not-blacklisted  1
v3-group-
1
      tun-1-l2tp-v3
-----
No. of tunnels: 1
=====

```

*A:Fden-Dut2-BSA2# show router l2tp peer 10.0.0.1 statistics

```

=====
Peer IP: 10.0.0.1
=====
tunnels           : 1
tunnels active    : 1
sessions          : 1
sessions active   : 1

rx ctrl octets    : 541
rx ctrl packets   : 5
tx ctrl octets    : 272
tx ctrl packets   : 5
tx error packets  : 0
rx error packets  : 0
rx accepted msg   : 4
rx duplicate msg   : 0
rx out of window msg : 0

acceptedMsgType
  StartControlConnectionRequest : 1
  StartControlConnectionConnected : 1
  IncomingCallRequest : 1
  IncomingCallConnected : 1
  ZeroLengthBody : 1
originalTransmittedMsgType
  StartControlConnectionReply : 1
  IncomingCallReply : 1

```

```

ZeroLengthBody                                     : 3
last cleared time                                  : N/A
=====

```

session

| | |
|--------------------|--|
| Syntax | session connection-id connection-id [detail] session [detail] [session-id session-id (v2)] [state session-state] [peer ip-address] [group group-name] [assignment-id assignment-id] [local-name local-host-name] [remote-name remote-host-name] [tunnel-id tunnel-id (v2)] session [detail] [state session-state] [peer ip-address] [group group-name] [assignment-id assignment-id] [local-name local-host-name] [remote-name remote-host-name] [control-connection-id connection-id (v3)] |
| Context | show>router>l2tp |
| Description | This command displays L2TP session operational information. |
| Parameters | connection-id connection-id — specifies the identification number for a Layer Two Tunneling Protocol connection Values 1 to 429496729 detail — displays detailed L2TP session information session-id session-id (v2) — specifies the identification number for a Layer Two Tunneling Protocol session Values 1 to 65535 state session-state — specifies the values to identify the operational state of the L2TP session Values closed, closed-by-peer, established, idle, wait-reply, wait-tunnel peer ip-address — specifies the IP address of the peer Values The following values apply to the 7750 SR: ipv4-address a.b.c.d (host bits must be 0) ipv6-address x:x:x:x:x:x:x[-interface] x:x:x:x:x:d.d.d.d[-interface] x: [0 to FFFF]H d: [0 to 255]D interface: 32 characters maximum, mandatory for link local addresses Values The following values apply to the 7450 ESS: ipv4-address: a.b.c.d (host bits must be 0) |

group *group-name* — specifies a string to identify a Layer Two Tunneling Protocol Tunnel group

assignment-id *assignment-id* — specifies a string that distinguishes this Layer Two Tunneling Protocol tunnel

local-name *local-host-name* — specifies the host name used by this system during the authentication phase of tunnel establishment

remote-name *remote-host-name* — specifies a string that is compared to the host name used by the tunnel peer during the authentication phase of tunnel establishment

tunnel-id *tunnel-id (v2)* — specifies the local identifier of this Layer Two Tunneling Protocol tunnel, when L2TP version 2 is used

Values 1 to 65535

control-connection-id *connection-id (v3)* — specifies an identification number for a Layer Two Tunneling Protocol session

Values 1 to 429496729

Output The following output is an example of L2TP session operational information.

Sample Output

```
*A:Dut-C# show router l2tp session
=====
L2TP Session Summary
=====
```

| ID | Control Conn ID | Tunnel-ID | Session-ID | State |
|-----------|-----------------|-----------|------------|-------------|
| 143524786 | 143523840 | 2190 | 946 | established |
| 143526923 | 143523840 | 2190 | 3083 | established |
| 143531662 | 143523840 | 2190 | 7822 | closed |
| 236926987 | 236912640 | 3615 | 14347 | closed |
| 236927915 | 236912640 | 3615 | 15275 | closed |
| 379407426 | 379387904 | 5789 | 19522 | established |
| 658187773 | 658178048 | 10043 | 9725 | established |
| 658198275 | 658178048 | 10043 | 20227 | established |
| 658210606 | 658178048 | 10043 | 32558 | established |

```
-----
No. of sessions: 9
=====
*A:Dut-C#

*A:Dut-C# show router l2tp session state established
=====
L2TP Session Summary
=====
```

| ID | Control Conn ID | Tunnel-ID | Session-ID | State |
|-----------|-----------------|-----------|------------|-------------|
| 143524786 | 143523840 | 2190 | 946 | established |
| 143526923 | 143523840 | 2190 | 3083 | established |
| 379407426 | 379387904 | 5789 | 19522 | established |
| 658187773 | 658178048 | 10043 | 9725 | established |
| 658198275 | 658178048 | 10043 | 20227 | established |


```

658210606          658178048          10043          32558          established
-----
No. of sessions: 6
=====
*A:Dut-C#

*A:Dut-C# show router l2tp session state closed detail
=====
L2TP Session Status
=====
Connection ID : 143531662
State          : closed
Tunnel Group   : ispl.group-2
Assignment ID  : ispl.tunnel-3
Error Message  : Terminated by PPPoE: RX PADT

Control Conn ID : 143523840          Remote Conn ID : 1148557524
Tunnel ID       : 2190              Remote Tunnel ID : 17525
Session ID      : 7822              Remote Session ID : 39124
Time Started    : 04/17/2009 18:44:37
Time Established : 04/17/2009 18:44:37 Time Closed      : 04/17/2009 18:44:50
CDN Result      : generalError       General Error    : noError
-----
L2TP Session Status
=====
Connection ID : 236926987
State          : closed
Tunnel Group   : ispl.group-2
Assignment ID  : ispl.tunnel-2
Error Message  : tunnel was closed

Control Conn ID : 236912640          Remote Conn ID : 3861360381
Tunnel ID       : 3615              Remote Tunnel ID : 58919
Session ID      : 14347             Remote Session ID : 44797
Time Started    : 04/17/2009 18:41:55
Time Established : 04/17/2009 18:41:55 Time Closed      : 04/17/2009 18:43:20
CDN Result      : generalError       General Error    : noError
-----
L2TP Session Status
=====
Connection ID : 236927915
State          : closed
Tunnel Group   : ispl.group-2
Assignment ID  : ispl.tunnel-2
Error Message  : tunnel was closed

Control Conn ID : 236912640          Remote Conn ID : 3861317210
Tunnel ID       : 3615              Remote Tunnel ID : 58919
Session ID      : 15275             Remote Session ID : 1626
Time Started    : 04/17/2009 18:41:03
Time Established : 04/17/2009 18:41:03 Time Closed      : 04/17/2009 18:43:20
CDN Result      : generalError       General Error    : noError
-----
No. of sessions: 3
=====
*A:Dut-C#

```

```

*A:Dut-C# show router l2tp session session-id 946
=====
L2TP Session Summary
=====
ID                Control Conn ID    Tunnel-ID    Session-ID    State
-----
143524786         143523840        2190         946           established
-----
No. of sessions: 1
=====
*A:Dut-C# show router l2tp session connection-id 143524786 detail
=====
L2TP Session Status
=====
Connection ID : 143524786
State          : established
Tunnel Group   : ispl.group-2
Assignment ID  : ispl.tunnel-3
Error Message  : N/A

Control Conn ID : 143523840      Remote Conn ID : 1148528691
Tunnel ID       : 2190          Remote Tunnel ID : 17525
Session ID      : 946           Remote Session ID : 10291
Time Started    : 04/17/2009 18:42:01
Time Established : 04/17/2009 18:42:01 Time Closed      : N/A
CDN Result      : noError       General Error    : noError
-----
*A:Dut-C#

*A:Dut-C# show router l2tp session group ispl.group-2
=====
L2TP Session Summary
=====
ID                Control Conn ID    Tunnel-ID    Session-ID    State
-----
143524786         143523840        2190         946           established
143526923         143523840        2190         3083          established
143531662         143523840        2190         7822          closed
236926987         236912640        3615         14347         closed
236927915         236912640        3615         15275         closed
658187773         658178048        10043        9725          established
658198275         658178048        10043        20227         established
658210606         658178048        10043        32558         established
-----
No. of sessions: 8
=====
*A:Dut-C#

*A:Dut-C# show router l2tp session tunnel-id 2190 state closed detail
=====
L2TP Session Status
=====
Connection ID : 143531662
State          : closed
Tunnel Group   : ispl.group-2

```

```
Assignment ID : ispl.tunnel-3
Error Message : Terminated by PPPoE: RX PADT

Control Conn ID   : 143523840           Remote Conn ID   : 1148557524
Tunnel ID         : 2190                 Remote Tunnel ID  : 17525
Session ID        : 7822                 Remote Session ID : 39124
Time Started      : 04/17/2009 18:44:37
Time Established  : 04/17/2009 18:44:37 Time Closed      : 04/17/2009 18:44:50
CDN Result        : generalError         General Error    : noError
```

No. of sessions: 1

=====

```
*A:Dut-C#
```

```
*A:Dut-C# show router l2tp session assignment-id ispl.tunnel-2
```

=====

```
L2TP Session Summary
```

| ID | Control Conn ID | Tunnel-ID | Session-ID | State |
|-----------|-----------------|-----------|------------|-------------|
| 236926987 | 236912640 | 3615 | 14347 | closed |
| 236927915 | 236912640 | 3615 | 15275 | closed |
| 658187773 | 658178048 | 10043 | 9725 | established |
| 658198275 | 658178048 | 10043 | 20227 | established |
| 658210606 | 658178048 | 10043 | 32558 | established |

No. of sessions: 5

=====

```
*A:Dut-C#
```

```
*A:Dut-C# show router l2tp session assignment-id ispl.tunnel-2 state established
```

=====

```
L2TP Session Summary
```

| ID | Control Conn ID | Tunnel-ID | Session-ID | State |
|-----------|-----------------|-----------|------------|-------------|
| 658187773 | 658178048 | 10043 | 9725 | established |
| 658198275 | 658178048 | 10043 | 20227 | established |
| 658210606 | 658178048 | 10043 | 32558 | established |

No. of sessions: 3

=====

```
*A:Dut-C#
```

```
*A:Dut-C# show router l2tp session control-connection-id 658178048
```

=====

```
L2TP Session Summary
```

| ID | Control Conn ID | Tunnel-ID | Session-ID | State |
|-----------|-----------------|-----------|------------|-------------|
| 658187773 | 658178048 | 10043 | 9725 | established |
| 658198275 | 658178048 | 10043 | 20227 | established |
| 658210606 | 658178048 | 10043 | 32558 | established |

No. of sessions: 3

=====

```
*A:Dut-C#
```

```
*A:Dut-C# show router l2tp session peer 10.10.20.100
```

```
=====
L2TP Session Summary
=====
```

| ID | Control Conn ID | Tunnel-ID | Session-ID | State |
|-----------|-----------------|-----------|------------|-------------|
| 236926987 | 236912640 | 3615 | 14347 | closed |
| 236927915 | 236912640 | 3615 | 15275 | closed |
| 658187773 | 658178048 | 10043 | 9725 | established |
| 658198275 | 658178048 | 10043 | 20227 | established |
| 658210606 | 658178048 | 10043 | 32558 | established |

```
-----
No. of sessions: 5
=====
```

```
*A:Dut-C#
```

```
*A:Dut-C# show router l2tp session peer 10.10.20.100 state closed detail
```

```
=====
L2TP Session Status
=====
```

```
Connection ID : 236926987
State          : closed
Tunnel Group   : ispl.group-2
Assignment ID  : ispl.tunnel-2
Error Message  : tunnel was closed
```

| | | | |
|------------------|-----------------------|-------------------|-----------------------|
| Control Conn ID | : 236912640 | Remote Conn ID | : 3861360381 |
| Tunnel ID | : 3615 | Remote Tunnel ID | : 58919 |
| Session ID | : 14347 | Remote Session ID | : 44797 |
| Time Started | : 04/17/2009 18:41:55 | | |
| Time Established | : 04/17/2009 18:41:55 | Time Closed | : 04/17/2009 18:43:20 |
| CDN Result | : generalError | General Error | : noError |

```
-----
L2TP Session Status
=====
```

```
Connection ID : 236927915
State          : closed
Tunnel Group   : ispl.group-2
Assignment ID  : ispl.tunnel-2
Error Message  : tunnel was closed
```

| | | | |
|------------------|-----------------------|-------------------|-----------------------|
| Control Conn ID | : 236912640 | Remote Conn ID | : 3861317210 |
| Tunnel ID | : 3615 | Remote Tunnel ID | : 58919 |
| Session ID | : 15275 | Remote Session ID | : 1626 |
| Time Started | : 04/17/2009 18:41:03 | | |
| Time Established | : 04/17/2009 18:41:03 | Time Closed | : 04/17/2009 18:43:20 |
| CDN Result | : generalError | General Error | : noError |

```
-----
No. of sessions: 2
=====
```

```
*A:Dut-C#
```

```
*A:Dut-C# show router l2tp session local-name lac1.wholesaler.com
```

```
=====
L2TP Session Summary
=====
ID                Control Conn ID    Tunnel-ID    Session-ID    State
-----
143524786         143523840        2190         946           established
143526923         143523840        2190         3083          established
143531662         143523840        2190         7822          closed
236926987         236912640        3615         14347         closed
236927915         236912640        3615         15275         closed
379407426         379387904        5789         19522         established
658187773         658178048        10043        9725          established
658198275         658178048        10043        20227         established
658210606         658178048        10043        32558         established
-----
No. of sessions: 9
=====
*A:Dut-C#
```

```
*A:Dut-C# show router l2tp session local-name lac1.wholesaler.com remote-
name lns.retailer1.net
=====
L2TP Session Summary
=====
ID                Control Conn ID    Tunnel-ID    Session-ID    State
-----
379407426         379387904        5789         19522         established
-----
No. of sessions: 1
=====
*A:Dut-C#
```

```
*A:Fden-Dut2-BSA2# show router l2tp session connection-id 600407016
=====
L2TP Session Summary
=====
ID                Control Conn ID    Tunnel-ID    Session-ID    State
-----
600407016         600375296        9161         31720         established
simon@base.lac.base.lns
interface: gi_base_lns_base_lac
service-id: 100
ip-address: 10.100.2.1
=====
```

```
*A:Fden-Dut2-BSA2# show router l2tp session connection-id 600407016 detail
=====
L2TP Session Status
=====

Connection ID: 600407016
State          : established
Tunnel Group  : base_lns_base_lac
Assignment ID : t1
Error Message  : N/A
```

```

Control Conn ID   : 600375296          Remote Conn ID    : 1026712216
Tunnel ID        : 9161                Remote Tunnel ID   : 15666
Session ID       : 31720                Remote Session ID  : 25240
Time Started     : 02/02/2010 09:08:54
Time Established : 02/02/2010 09:08:54 Time Closed       : N/A
CDN Result       : noError              General Error     : noError
-----

```

PPP information

```

Service Id       : 100
Interface        : gi_base_lns_base_lac
LCP State        : opened
IPCP State       : opened
IPv6CP State     : initial
PPP MTU          : 1492
PPP Auth-Protocol : chap
PPP User-Name    : simon@base.lac.base.lns

```

```

Subscriber Origin : radius
Strings Origin    : radius
IPCP Info Origin  : radius
IPv6CP Info Origin : none

```

```

Subscriber       : "simon"
Sub-Profile-String : "sub1"
SLA-Profile-String : "sla1"
ANCP-String      : ""
Int-Dest-Id      : ""
App-Profile-String : ""
Category-Map-Name : ""

```

```

IP Address       : 10.100.2.1
Primary DNS      : N/A
Secondary DNS    : N/A
Primary NBNS     : N/A
Secondary NBNS   : N/A
Address-Pool     : N/A

```

```

IPv6 Prefix      : N/A
IPv6 Del.Pfx.    : N/A
Primary IPv6 DNS  : N/A
Secondary IPv6 DNS : N/A

```

```

Circuit-Id       : (Not Specified)
Remote-Id        : (Not Specified)

```

```

Session-Timeout  : N/A
Radius Class     : (Not Specified)
Radius User-Name  : simon@base.lac.base.lns

```

statistics

Syntax **statistics****Context** **show>router>l2tp**

Description This command displays L2TP statistics.

Output The following output is an example of L2TP statistics information.

Sample Output

```
*A:Dut-C# show router l2tp statistics
=====
L2TP Statistics
=====
Tunnels                               Sessions
-----
Active           : 3                   Active           : 6

Setup history since 04/17/2009 18:38:41

Total           : 4                   Total           : 9
Failed          : 0                   Failed          : 0
Failed Auth     : 0
=====
*A:Dut-C#
```

tunnel

Syntax **tunnel** [**statistics**] [**detail**] [**peer** *ip-address*] [**state** *tunnel-state*] [**remote-connection-id** *remote-connection-id (v3)*] [**group** *group-name*] [**assignment-id** *assignment-id*] [**local-name** *host-name*] [**remote-name** *host-name*]
tunnel [**statistics**] [**detail**] [**peer** *ip-address*] [**state** *tunnel-state*] [**remote-tunnel-id** *remote-tunnel-id (v2)*] [**group** *group-name*] [**assignment-id** *assignment-id*] [**local-name** *host-name*] [**remote-name** *host-name*]
tunnel *tunnel-id* *tunnel-id (v2)* [**statistics**] [**detail**]
tunnel *connection-id* *connection-id (v3)* [**statistics**] [**detail**]

Context show>router>l2tp

Description This command displays L2TP tunnel operational information.

Parameters **statistics** — displays L2TP tunnel statistics
detail — displays detailed L2TP tunnel information
peer *ip-address* — displays information for the IP address of the peer
state *tunnel-state* — displays the operational state of the tunnel
remote-connection-id *remote-connection-id (v3)* — displays information for the specified remote connection ID
group *group-name* — displays L2TP tunnel information for the specified tunnel group
assignment-id *assignment-id* — specifies a string that distinguishes this Layer Two Tunneling Protocol tunnel
local-name *host-name* — specifies a local host name used by this system

remote-name *host-name* — specifies a remote host name used by this system

connection-id *connection-id* — specifies the identification number for a Layer Two Tunneling Protocol connection

Values 1 to 429496729

detail — displays detailed L2TP session information

session-id *session-id* (v2) — displays information for the specified the L2TP session

Values 1 to 65535

state *session-state* — displays the operational state of the L2TP session

Values closed, closed-by-peer, draining, drained, established, established-idle, idle, wait-reply, wait-conn

peer *ip-address* — displays information for the specified peer IP address

Values The following values apply to the 7750 SR:

ipv4-address a.b.c.d (host bits must be 0)
 ipv6-address x:x:x:x:x:x:x[-interface]
 x:x:x:x:x:x:d.d.d.d[-interface]
 x: [0 to FFFF]H
 d: [0 to 255]D
 interface: 32 characters maximum, mandatory for link local addresses

Values The following values apply to the 7450 ESS:

ipv4-address: a.b.c.d (host bits must be 0)

tunnel-id *tunnel-id* (v2) — displays information for the specified ID of a L2TP tunnel.

In L2TP version 2, it is the 16-bit tunnel ID

Values 1 to 65535

control-connection-id *connection-id* (v3) — displays information for the specified ID of a L2TP tunnel. In L2TP version 3, it is the 32-bit control connection ID

Values 1 to 429496729

Output The following output is an example of L2TP tunnel operational information.

Sample Output

```
*A:Dut-C# show router l2tp tunnel
=====
Conn ID          Loc-Tu-ID Rem-Tu-ID State          Ses Active
Group           Assignment          Ses Total
-----
143523840        2190      17525   established        2
isp1.group-2          3
```



```

isp1.tunnel-3
236912640          3615      58919      closedByPeer      0
isp1.group-2
isp1.tunnel-2
379387904          5789      4233      established        1
isp1.group-1
isp1.tunnel-1
658178048          10043     33762     draining          3
isp1.group-2
isp1.tunnel-2

```

No. of tunnels: 4

*A:Dut-C#

*A:Dut-C# show router l2tp tunnel state closed-by-peer detail

L2TP Tunnel Status

```

Connection ID : 236912640
State          : closedByPeer
IP             : 10.20.1.3
Peer IP        : 10.10.20.100
Name           : lacl.wholesaler.com
Remote Name    : lns2.retailer1.net
Assignment ID  : isp1.tunnel-2
Group Name     : isp1.group-2
Error Message  : Goodbye!

```

```

Tunnel ID      : 3615
UDP Port       : 1701
Preference     : 100
Hello Interval (s) : infinite
Idle TO (s)    : 60
Max Retr Estab : 5
Session Limit  : 1000
Transport Type : udpIp
Time Started   : 04/17/2009 18:41:03
Time Established : 04/17/2009 18:41:03
Stop CCN Result : generalReq

Remote Conn ID : 3861315584
Remote Tunnel ID : 58919
Remote UDP Port : 1701
Destruct TO (s) : 7200
Max Retr Not Estab: 5
AVP Hiding      : never
Challenge       : never
Time Idle       : 04/17/2009 18:43:20
Time Closed     : 04/17/2009 18:43:20
General Error   : noError

```

No. of tunnels: 1

*A:Dut-C#

*A:Dut-C# show router l2tp tunnel state established

```

Conn ID          Loc-Tu-ID Rem-Tu-ID State      Ses Active
Group            Assignment
-----
143523840        2190      17525     established  2
isp1.group-2
isp1.tunnel-3
379387904        5789      4233     established  1
isp1.group-1

```

```
isp1.tunnel-1
-----
No. of tunnels: 2
=====
*A:Dut-C#

*A:Dut-C# show router l2tp tunnel tunnel-id 2190 statistics
=====
L2TP Tunnel Statistics
=====
Connection ID: 143523840
-----

```

| | Attempts | Failed | Active | Total |
|----------|----------|--------|--------|-------|
| Sessions | 3 | 0 | 2 | 3 |

```
-----
Rx                                     Tx
-----
Ctrl Packets  47                       47
Ctrl Octets   954                      1438
Error Packets 0                        0
-----
*A:Dut-C#

*A:Dut-C# show router l2tp tunnel connection-id 143523840 statistics
=====
L2TP Tunnel Statistics
=====
Connection ID: 143523840
-----

```

| | Attempts | Failed | Active | Total |
|----------|----------|--------|--------|-------|
| Sessions | 3 | 0 | 2 | 3 |

```
-----
Rx                                     Tx
-----
Ctrl Packets  48                       48
Ctrl Octets   974                      1450
Error Packets 0                        0
-----
*A:Dut-C#

*A:Dut-C# show router l2tp tunnel remote-tunnel-id 17525 detail
=====
L2TP Tunnel Status
=====
Connection ID : 143523840
State         : established
IP            : 10.20.1.3
Peer IP       : 10.10.20.101
Name          : lacl.wholesaler.com
Remote Name   : lns3.retailer1.net
Assignment ID : isp1.tunnel-3
Group Name    : isp1.group-2
```

Error Message : N/A

| | | | |
|--------------------|-----------------------|--------------------|--------------|
| Tunnel ID | : 2190 | Remote Conn ID | : 1148518400 |
| UDP Port | : 1701 | Remote Tunnel ID | : 17525 |
| Preference | : 100 | Remote UDP Port | : 1701 |
| Hello Interval (s) | : 300 | | |
| Idle TO (s) | : 0 | Destruct TO (s) | : 7200 |
| Max Retr Estab | : 5 | Max Retr Not Estab | : 5 |
| Session Limit | : 1000 | AVP Hiding | : never |
| Transport Type | : udpIp | Challenge | : never |
| Time Started | : 04/17/2009 18:41:14 | Time Idle | : N/A |
| Time Established | : 04/17/2009 18:41:14 | Time Closed | : N/A |
| Stop CCN Result | : noError | General Error | : noError |

No. of tunnels: 1

=====

*A:Dut-C# show router l2tp tunnel remote-connection-id 1148518400 statistics

=====

L2TP Tunnel Statistics

=====

Connection ID: 143523840

| | Attempts | Failed | Active | Total |
|---------------|----------|--------|--------|-------|
| ----- | ----- | ----- | ----- | ----- |
| Sessions | 3 | 0 | 2 | 3 |
| ----- | ----- | ----- | ----- | ----- |
| | Rx | | Tx | |
| ----- | ----- | ----- | ----- | ----- |
| Ctrl Packets | 50 | | 50 | |
| Ctrl Octets | 1014 | | 1474 | |
| Error Packets | 0 | | 0 | |
| ----- | ----- | ----- | ----- | ----- |

No. of tunnels: 1

=====

*A:Dut-C# show router l2tp tunnel peer 10.10.20.100 state closed-by-peer detail

=====

L2TP Tunnel Status

=====

Connection ID : 236912640
State : closedByPeer
IP : 10.20.1.3
Peer IP : 10.10.20.100
Name : lacl.wholesaler.com
Remote Name : lns2.retailer1.net
Assignment ID : ispl.tunnel-2
Group Name : ispl.group-2
Error Message : Goodbye!

| | | | |
|-----------|--------|------------------|--------------|
| Tunnel ID | : 3615 | Remote Conn ID | : 3861315584 |
| UDP Port | : 1701 | Remote Tunnel ID | : 58919 |
| | | Remote UDP Port | : 1701 |

```
Preference          : 100
Hello Interval (s)  : infinite
Idle TO (s)         : 60
Max Retr Estab      : 5
Session Limit       : 1000
Transport Type      : udpIp
Time Started        : 04/17/2009 18:41:03
Time Established    : 04/17/2009 18:41:03
Stop CCN Result     : generalReq
Destruct TO (s)     : 7200
Max Retr Not Estab  : 5
AVP Hiding          : never
Challenge           : never
Time Idle           : 04/17/2009 18:43:20
Time Closed         : 04/17/2009 18:43:20
General Error       : noError
```

No. of tunnels: 1

*A:Dut-C#

*A:Dut-C# show router l2tp tunnel group ispl.group-2

```
=====
Conn ID          Loc-Tu-ID Rem-Tu-ID State          Ses Active
Group                                                    Ses Total
Assignment
-----
143523840        2190      17525   established          2
  ispl.group-2                                     3
  ispl.tunnel-3
236912640        3615      58919   closedByPeer         0
  ispl.group-2                                     2
  ispl.tunnel-2
658178048        10043     33762   draining             3
  ispl.group-2                                     3
  ispl.tunnel-2
-----
```

No. of tunnels: 3

*A:Dut-C#

*A:Dut-C# show router l2tp tunnel assignment-id ispl.tunnel-3 state established statistics

L2TP Tunnel Statistics

Connection ID: 143523840

```
-----
Attempts  Failed          Active  Total
-----
Sessions   3           0           2       3
-----
Rx                                     Tx
-----
Ctrl Packets  66                                     66
Ctrl Octets   1310                                    1690
Error Packets 0                                     0
-----
```

No. of tunnels: 1

*A:Dut-C#

```
*A:Dut-C# show router l2tp tunnel local-name lac1.wholesaler.com remote-
name lns2.retailer1.net state draining
=====
Conn ID                      Loc-Tu-ID Rem-Tu-ID State              Ses Active
  Group                               Ses Total
  Assignment
-----
658178048                    10043     33762    draining              3
    ispl.group-2                               3
    ispl.tunnel-2
-----
No. of tunnels: 1
=====
*A:Dut-C#

*A:Fden-Dut2-BSA2# show router l2tp tunnel connection-id 600375296 statistics
=====
L2TP Tunnel Statistics
=====

Connection ID: 600375296

-----
Attempts    Failed              Active    Total
-----
Sessions    1          0              1         1
-----

-----
Rx              Tx
-----
Ctrl Packets   6              6
Ctrl Octets    553            292
Error Packets  0              0
-----

-----
Accepted    Duplicate          Out-Of-Wnd
-----
Fsm Messages 4          0              0
-----

-----
Unsent Max Unsent Cur          Ack Max    Ack Cur
-----
Q Length     1          0              1         0
-----

Window Size Cur          : 4
acceptedMsgType
  StartControlConnectionRequest      : 1
  StartControlConnectionConnected    : 1
  IncomingCallRequest                 : 1
  IncomingCallConnected               : 1
  ZeroLengthBody                      : 3
originalTransmittedMsgType
  StartControlConnectionReply         : 1
  Hello                              : 2
```

```

IncomingCallReply          : 1
ZeroLengthBody             : 3

last cleared time          : N/A
=====

```

On LAC (master node after switchover)

```

=====
L2TP Tunnel Status
=====

```

```

Connection ID: 11206656
State          : established
IP             : 10.124.0.9
UDP            : 1701
Peer IP        : 10.124.0.3
Peer UDP       : 1701
Tx dst-IP      : 10.124.0.3
Tx dst-UDP     : 1701
Rx src-IP      : 10.124.0.3
Rx src-UDP     : 1701
Name           : mc-lac
Remote Name    : mc-lns
Assignment ID: t1
Group Name     : mc-lac
Acct. Policy   : l2tp-base
Error Message: N/A

```

```

Tunnel ID          : 171
Preference         : 50
Hello Interval (s) : infinite
Idle TO (s)        : infinite
Max Retr Estab     : 5
Session Limit      : 32767
Transport Type     : udpIp
Time Started       : 02/19/2015 13:00:36
Time Established   : 02/19/2015 13:00:36
Stop CCN Result    : noError
Blacklist-state    : not-blacklisted
Set Dont Fragment : true

Remote Conn ID     : 429260800
Remote Tunnel ID   : 6550
Receive Window     : 64
Destruct TO (s)    : 60
Max Retr Not Estab : 5
AVP Hiding         : never
Challenge          : never
Time Idle          : N/A
Time Closed        : N/A
General Error      : noError

```

```

Failover
State          : recoverable
Recovery Conn ID : N/A
Recovery state   : not-applicable
Recovered Conn ID : N/A
Recovery method  : mcs
Track SRRP      : 124
Ctrl msg behavior : handle

```

```

-----
No. of tunnels: 1
=====

```

On LAC (slave node after switchover)

show router l2tp tunnel detail

=====

L2TP Tunnel Status

=====

Connection ID: 11206656
 State : draining
 IP : 10.124.0.9
 UDP : 1701
 Peer IP : 10.124.0.3
 Peer UDP : 1701
 Tx dst-IP : 10.124.0.3
 Tx dst-UDP : 1701
 Rx src-IP : 10.124.0.3
 Rx src-UDP : 1701
 Name : mc-lac
 Remote Name : mc-lns
 Assignment ID: t1
 Group Name : mc-lac
 Acct. Policy : l2tp-base
 Error Message: N/A

| | | | |
|--------------------|-----------------------|--------------------|-------------|
| Tunnel ID | : 171 | Remote Conn ID | : 429260800 |
| Preference | : 50 | Remote Tunnel ID | : 6550 |
| Hello Interval (s) | : infinite | Receive Window | : 64 |
| Idle TO (s) | : infinite | Destruct TO (s) | : 60 |
| Max Retr Estab | : 5 | Max Retr Not Estab | : 5 |
| Session Limit | : 32767 | AVP Hiding | : never |
| Transport Type | : udpIp | Challenge | : never |
| Time Started | : 02/19/2015 13:00:36 | Time Idle | : N/A |
| Time Established | : 02/19/2015 13:00:36 | Time Closed | : N/A |
| Stop CCN Result | : noError | General Error | : noError |
| Blacklist-state | : not-blacklisted | | |
| Set Dont Fragment | : true | | |

Failover

State : recoverable

Recovery Conn ID : N/A

Recovery state : not-applicable

Recovered Conn ID : N/A

Recovery method : mcs

Track SRRP : 124

Ctrl msg behavior : forward-to-mcs-peer

No. of tunnels: 1

=====

On LNS after switchover

show router l2tp tunnel detail

=====

L2TP Tunnel Status

=====

Connection ID: 429260800
 State : established

```

IP          : 10.124.0.3
UDP         : 1701
Peer IP     : 10.124.0.9
Peer UDP    : 1701
Tx dst-IP   : 10.124.0.9
Tx dst-UDP  : 1701
Rx src-IP   : 10.124.0.9
Rx src-UDP  : 1701
Name        : mc-lns
Remote Name : mc-lac
Assignment ID: t1
Group Name  : mc-lns
Acct. Policy : N/A
Error Message: N/A

```

```

Tunnel ID      : 6550
Preference     : 50
Hello Interval (s) : 300
Idle TO (s)    : infinite
Max Retr Estab : 5
Session Limit  : 32767
Transport Type : udpIp
Time Started   : 02/19/2015 13:00:36
Time Established : 02/19/2015 13:00:36
Stop CCN Result : noError
Blacklist-state : not-blacklisted
Set Dont Fragment : true

Remote Conn ID : 11206656
Remote Tunnel ID : 171
Receive Window  : 64
Destruct TO (s) : 60
Max Retr Not Estab : 5
AVP Hiding      : never
Challenge       : never
Time Idle       : N/A
Time Closed     : N/A
General Error   : noError

```

```

Failover
State      : not-recoverable
Recovery Conn ID : N/A
Recovery state : not-applicable
Recovered Conn ID : N/A
Recovery method : mcs
Track SRRP    : (Not specified)
Ctrl msg behavior : handle

```

```

-----
No. of tunnels: 1
=====

```

On LAC (master node after switchover; 7536640 is the recovered tunnel, 1865089024 is the recovery tunnel)

```

=====
L2TP Tunnel Status
=====
Connection ID: 7536640
State      : established
IP         : 10.124.0.9
UDP        : 1701
Peer IP    : 10.124.0.3
Peer UDP   : 1701
Tx dst-IP  : 10.124.0.3
Tx dst-UDP : 1701
Rx src-IP  : 10.124.0.3
Rx src-UDP : 1701
Name       : mc-lac
Remote Name : mc-lns

```


Assignment ID: t1
Group Name : mc-lac
Acct. Policy : l2tp-base
Error Message: N/A

| | | | |
|--------------------|-----------------------|--------------------|-------------|
| Tunnel ID | : 115 | Remote Conn ID | : 433324032 |
| Preference | : 50 | Remote Tunnel ID | : 6612 |
| Hello Interval (s) | : infinite | Receive Window | : 64 |
| Idle TO (s) | : infinite | Destruct TO (s) | : 60 |
| Max Retr Estab | : 5 | Max Retr Not Estab | : 5 |
| Session Limit | : 32767 | AVP Hiding | : never |
| Transport Type | : udpIp | Challenge | : never |
| Time Started | : 02/19/2015 13:07:53 | Time Idle | : N/A |
| Time Established | : 02/19/2015 13:07:53 | Time Closed | : N/A |
| Stop CCN Result | : noError | General Error | : noError |
| Blacklist-state | : not-blacklisted | | |
| Set Dont Fragment | : true | | |

Failover
State : recoverable
Recovery Conn ID : 1865089024
Recovery state : not-applicable
Recovered Conn ID : N/A
Recovery method : recovery-tunnel
Track SRRP : 124
Ctrl msg behavior : handle

Connection ID: 1865089024
State : closed
IP : 10.124.0.9
UDP : 1701
Peer IP : 10.124.0.3
Peer UDP : 1701
Tx dst-IP : 10.124.0.3
Tx dst-UDP : 1701
Rx src-IP : 10.124.0.3
Rx src-UDP : 1701
Name : mc-lac
Remote Name : mc-lns
Assignment ID: t1
Group Name : mc-lac
Acct. Policy : l2tp-base
Error Message: N/A

| | | | |
|--------------------|-----------------------|--------------------|-----------------------|
| Tunnel ID | : 28459 | Remote Conn ID | : 1169424384 |
| Preference | : 50 | Remote Tunnel ID | : 17844 |
| Hello Interval (s) | : infinite | Receive Window | : 64 |
| Idle TO (s) | : 60 | Destruct TO (s) | : 60 |
| Max Retr Estab | : 5 | Max Retr Not Estab | : 5 |
| Session Limit | : 32767 | AVP Hiding | : never |
| Transport Type | : udpIp | Challenge | : never |
| Time Started | : 02/19/2015 13:12:05 | Time Idle | : N/A |
| Time Established | : 02/19/2015 13:12:05 | Time Closed | : 02/19/2015 13:12:05 |
| Stop CCN Result | : generalReq | General Error | : noError |
| Blacklist-state | : not-blacklisted | | |

Set Dont Fragment : true

Failover
State : not-applicable
Recovery Conn ID : N/A
Recovery state : recovery-tunnel
Recovered Conn ID : 7536640
Recovery method : default
Track SRRP : 124
Ctrl msg behavior : handle

No. of tunnels: 2
=====

On LAC (slave node after switchover)

=====

L2TP Tunnel Status

=====

Connection ID: 7536640
State : draining
IP : 10.124.0.9
UDP : 1701
Peer IP : 10.124.0.3
Peer UDP : 1701
Tx dst-IP : 10.124.0.3
Tx dst-UDP : 1701
Rx src-IP : 10.124.0.3
Rx src-UDP : 1701
Name : mc-lac
Remote Name : mc-lns
Assignment ID: t1
Group Name : mc-lac
Acct. Policy : l2tp-base
Error Message: N/A

| | | | |
|--------------------|-----------------------|--------------------|-------------|
| Tunnel ID | : 115 | Remote Conn ID | : 433324032 |
| Preference | : 50 | Remote Tunnel ID | : 6612 |
| Hello Interval (s) | : infinite | Receive Window | : 64 |
| Idle TO (s) | : infinite | Destruct TO (s) | : 60 |
| Max Retr Estab | : 5 | Max Retr Not Estab | : 5 |
| Session Limit | : 32767 | AVP Hiding | : never |
| Transport Type | : udpIp | Challenge | : never |
| Time Started | : 02/19/2015 13:07:53 | Time Idle | : N/A |
| Time Established | : 02/19/2015 13:07:53 | Time Closed | : N/A |
| Stop CCN Result | : noError | General Error | : noError |
| Blacklist-state | : not-blacklisted | | |
| Set Dont Fragment | : true | | |

Failover
State : recoverable
Recovery Conn ID : N/A
Recovery state : not-applicable
Recovered Conn ID : N/A
Recovery method : recovery-tunnel
Track SRRP : 124
Ctrl msg behavior : forward-to-mcs-peer

No. of tunnels: 1
=====

On LNS after switchover (433324032 is the recovered tunnel, 1169424384 is the recovery tunnel)

=====

L2TP Tunnel Status

=====

Connection ID: 433324032
State : established
IP : 10.124.0.3
UDP : 1701
Peer IP : 10.124.0.9
Peer UDP : 1701
Tx dst-IP : 10.124.0.9
Tx dst-UDP : 1701
Rx src-IP : 10.124.0.9
Rx src-UDP : 1701
Name : mc-lns
Remote Name : mc-lac
Assignment ID: t1
Group Name : mc-lns
Acct. Policy : N/A
Error Message: N/A

| | | | |
|--------------------|-----------------------|--------------------|-----------|
| Tunnel ID | : 6612 | Remote Conn ID | : 7536640 |
| Preference | : 50 | Remote Tunnel ID | : 115 |
| Hello Interval (s) | : 300 | Receive Window | : 64 |
| Idle TO (s) | : infinite | Destruct TO (s) | : 60 |
| Max Retr Estab | : 5 | Max Retr Not Estab | : 5 |
| Session Limit | : 32767 | AVP Hiding | : never |
| Transport Type | : udpIp | Challenge | : never |
| Time Started | : 02/19/2015 13:07:53 | Time Idle | : N/A |
| Time Established | : 02/19/2015 13:07:53 | Time Closed | : N/A |
| Stop CCN Result | : noError | General Error | : noError |
| Blacklist-state | : not-blacklisted | | |
| Set Dont Fragment | : true | | |

Failover
State : not-recoverable
Recovery Conn ID : 1169424384
Recovery state : not-applicable
Recovered Conn ID : N/A
Recovery method : recovery-tunnel
Track SRRP : (Not specified)
Ctrl msg behavior : handle

Connection ID: 1169424384
State : closed
IP : 10.124.0.3
UDP : 1701
Peer IP : 10.124.0.9
Peer UDP : 1701
Tx dst-IP : 10.124.0.9

```

Tx dst-UDP      : 1701
Rx src-IP       : 10.124.0.9
Rx src-UDP      : 1701
Name            : mc-lns
Remote Name     : mc-lac
Assignment ID: t1
Group Name      : mc-lns
Acct. Policy    : N/A
Error Message   : N/A

Tunnel ID       : 17844
Preference      : 50
Hello Interval (s): infinite
Idle TO (s)     : 60
Max Retr Estab  : 5
Session Limit   : 32767
Transport Type  : udpIp
Time Started    : 02/19/2015 13:12:05
Time Established : 02/19/2015 13:12:05
13:12:05
Stop CCN Result : generalReq
Blacklist-state : not-blacklisted
Set Dont Fragment : true

Remote Conn ID   : 1865089024
Remote Tunnel ID : 28459
Receive Window   : 64
Destruct TO (s)  : 60
Max Retr Not Estab: 5
AVP Hiding       : never
Challenge        : never
Time Idle        : N/A
Time Closed      : 02/19/2015
13:12:05
General Error    : noError

Failover
State            : not-applicable
Recovery Conn ID : N/A
Recovery state   : recovery-tunnel
Recovered Conn ID : 433324032
Recovery method  : default
Track SRRP       : (Not specified)
Ctrl msg behavior : handle
-----
No. of tunnels: 2
=====

```

2.13.2.2 Clear Commands

The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

router

| | |
|--------------------|---|
| Syntax | router [<i>router-instance</i>] router service-name <i>service-name</i> |
| Context | clear>router |
| Description | This command enters the context in which to clear various parameters for the specified <i>router-instance</i> . |

Parameters *router-instance* — specifies the router name, CPM router instance, or service ID

Values *router-name* or *service-id*

router-instance : *router-name*

router-name Base | management | vpls-management |
cpm-vr-name

cpm-vr-name [32 characters maximum]

service-id: 1 to 2147483647

Default Base

service-name — specifies the service name, up to 64 characters

arp

Syntax **arp** {**all** | *ip-addr* | **interface** {*ip-int-name* | *ip-addr*}}

Context clear>router

Description This command clears all or specific ARP entries.

The scope of ARP cache entries cleared depends on the command line options specified.

Parameters **all** — clears all ARP cache entries

ip-addr — clears the ARP cache entry for the specified IP address

interface *ip-int-name* — clears all ARP cache entries for the IP interface with the specified name

interface *ip-addr* — clears all ARP cache entries for the specified IP interface with the specified IP address

bfd

Syntax **bfd src-ip** *ip-address* **dst-ip** *ip-address*
bfd all

Context clear>router

Description This command enables the context to clear bidirectional forwarding (BFD) sessions and statistics.

session

| | |
|--------------------|---|
| Syntax | session src-ip <i>ip-address</i> dst-ip <i>ip-address</i> |
| Context | clear>router>bfd |
| Description | This command clears BFD sessions. |
| Parameters | src-ip <i>ip-address</i> — specifies the address of the local endpoint of this BFD session dst-ip <i>ip-address</i> — specifies the address of the remote endpoint of this BFD session |

statistics

| | |
|--------------------|--|
| Syntax | statistics src-ip <i>ip-address</i> dst-ip <i>ip-address</i> statistics all |
| Context | clear>router>bfd |
| Description | This command clears BFD statistics. |
| Parameters | src-ip <i>ip-address</i> — specifies the address of the local endpoint of this BFD session dst-ip <i>ip-address</i> — specifies the address of the remote endpoint of this BFD session all — clears statistics for all BFD sessions |

dhcp

| | |
|--------------------|---|
| Syntax | dhcp |
| Context | clear>router |
| Description | This command enables the context to clear DHCP related information. |

dhcp6

| | |
|--------------------|--|
| Syntax | dhcp6 |
| Context | clear>router |
| Description | This command enables the context to clear DHCP6 related information. |

forwarding-table

| | |
|---------------|--|
| Syntax | forwarding-table [<i>slot-number</i>] |
|---------------|--|

| | |
|--------------------|--|
| Context | clear>router |
| Description | This command clears entries in the forwarding table (maintained by the IOMs). If the slot number is not specified, the command forces the route table to be recalculated. |
| Parameters | <i>slot-number</i> — clears the specified card slot |
| | Default all IOMs or linecards |
| | Values 1 to 10 |

grt-lookup

| | |
|--------------------|---|
| Syntax | grt-lookup |
| Context | clear>router |
| Description | This command re-evaluates route policies for GRT. |

icmp

| | |
|--------------------|---|
| Syntax | icmp all icmp global icmp interface <i>interface-name</i> |
| Context | clear>router |
| Description | This command clears ICMP statistics. |
| Parameters | all — clears all statistics global — clears global router statistics <i>interface-name</i> — clears ICMP statistics for the specified interface |
| | Values 32 characters maximum |

icmp-redirect-route

| | |
|--------------------|--|
| Syntax | icmp-redirect-route { all <i>ip-address</i> } |
| Context | clear>router |
| Description | This command deletes routes created as a result of ICMP redirects received on the management interface. |
| Parameters | all — clears all routes <i>ip-address</i> — clears the routes associated with the specified IP address |

icmp6

| | |
|--------------------|---|
| Syntax | icmp6 all icmp6 global icmp6 interface <i>interface-name</i> |
| Context | clear>router |
| Description | This command clears ICMPv6 statistics. |
| Parameters | all — clears all statistics global — clears global router statistics <i>interface-name</i> — clears ICMPv6 statistics for the specified interface |

interface

| | |
|--------------------|---|
| Syntax | interface [<i>ip-int-name</i> <i>ip-address</i>] [urpf-stats] [statistics] [hold-time] interface [<i>ip-int-name</i> <i>ip-address</i>] policy-accounting [class] [index] interface <i>ip-int-name</i> <i>ip-address</i> mac [<i>ieee-address</i>] |
| Context | clear>router |
| Description | This command clears IP interface statistics. If no IP interface is specified either by IP interface name or IP address, the command will perform the clear operation on all IP interfaces. |
| Parameters | <i>ip-int-name</i> <i>ip-addr</i> — Specifies IP interface name or IP interface address. Values ip-int-name: 32 chars max ip-address: a.b.c.d Default all IP interfaces icmp — Specifies to reset the ICMP statistics for the IP interface(s) used for ICMP rate limit. urpf-stats — Resets the statistics associated with uRPF failures. statistics — Resets the IP interface traffic statistics. hold-time — Clears the IP interface activation hold time. policy-accounting — Clears the accounting statistics. class — Specifies whether to clear source class counters or destination class counters. <i>index</i> — Specifies the source-class or destination-class ID. Values 1 to 255 |

ieee-address — Specifies the MAC address.

Values *xx:xx:xx:xx:xx:xx* or *xx-xx-xx-xx-xx-xx*

l2tp

| | |
|--------------------|--|
| Syntax | l2tp |
| Context | clear>router |
| Description | This command enables the context to clear L2PT data. |

group

| | |
|--------------------|---|
| Syntax | group <i>tunnel-group-name</i> |
| Context | clear>router>l2tp |
| Description | This command clears L2PT data. |
| Parameters | <i>tunnel-group-name</i> — specifies a Layer Two Tunneling Protocol Tunnel Group name |

tunnel

| | |
|--------------------|--|
| Syntax | tunnel <i>tunnel-id</i> |
| Context | clear>router>l2tp |
| Description | This command clears L2PT data. |
| Parameters | <i>tunnel-group-name</i> — clears L2TP tunnel statistics |

statistics

| | |
|--------------------|---|
| Syntax | statistics |
| Context | clear>router>l2tp clear>router>l2tp>group clear>router>l2tp> tunnel |
| Description | This command clears statistics for the specified context. |

statistics

| | |
|--------------------|---|
| Syntax | statistics [<i>ip-address</i> <i>ip-int-name</i>] |
| Context | clear>router>dhcp clear>router>dhcp6 |
| Description | <p>This command clear statistics for DHCP and DHCP6and DHCP6 relay and snooping statistics.</p> <p>If no IP address or interface name is specified, then statistics are cleared for all configured interfaces.</p> <p>If an IP address or interface name is specified, then only data regarding the specified interface is cleared.</p> |
| Parameters | <i>ip-address</i> <i>ip-int-name</i> — displays statistics for the specified IP interface |

neighbor

| | | | | | | | | | |
|--------------------|--|---------------|-------------------------------------|--|-------------------|--|-----------------|--|----------------|
| Syntax | neighbor { all <i>ip-address</i> } neighbor [interface <i>ip-int-name</i> <i>ip-address</i>] | | | | | | | | |
| Context | clear>router | | | | | | | | |
| Description | This command clears IPv6 neighbor information. | | | | | | | | |
| Parameters | <p>all — clears IPv6 neighbors</p> <p><i>ip-int-name</i> — clears the specified neighbor interface information</p> <p>Values 32 characters maximum</p> <p><i>ip-address</i> — clears the specified IPv6 neighbors</p> <p>Values</p> <table><tr><td>ipv6-address:</td><td>x:x:x:x:x:x:x (eight 16-bit pieces)</td></tr><tr><td></td><td>x:x:x:x:x:d.d.d.d</td></tr><tr><td></td><td>x: [0 to FFFF]H</td></tr><tr><td></td><td>d: [0 to 255]D</td></tr></table> | ipv6-address: | x:x:x:x:x:x:x (eight 16-bit pieces) | | x:x:x:x:x:d.d.d.d | | x: [0 to FFFF]H | | d: [0 to 255]D |
| ipv6-address: | x:x:x:x:x:x:x (eight 16-bit pieces) | | | | | | | | |
| | x:x:x:x:x:d.d.d.d | | | | | | | | |
| | x: [0 to FFFF]H | | | | | | | | |
| | d: [0 to 255]D | | | | | | | | |

router-advertisement

| | |
|----------------|---|
| Syntax | router-advertisement all router-advertisement [interface <i>interface-name</i>] |
| Context | clear>router |

| | |
|--------------------|---|
| Description | This command clears all router advertisement counters. |
| Parameters | <p><i>all</i> — clears all router advertisement counters for all interfaces</p> <p>interface <i>interface-name</i> — clear router advertisement counters for the specified interface</p> |

2.13.2.3 Debug Commands

destination

| | |
|--------------------|---|
| Syntax | destination <i>trace-destination</i> |
| Context | debug>trace |
| Description | This command specifies the destination to send trace messages. |
| Parameters | <i>trace-destination</i> — the destination to send trace messages |
| Values | stdout, console, logger, memory |

enable

| | |
|--------------------|---|
| Syntax | [no] enable |
| Context | debug>trace |
| Description | <p>This command enables the trace.</p> <p>The no form of the command disables the trace.</p> |

trace-point

| | |
|--------------------|--|
| Syntax | [no] trace-point [module <i>module-name</i>] [type <i>event-type</i>] [class <i>event-class</i>] [task <i>task-name</i>] [function <i>function-name</i>] |
| Context | debug>trace |
| Description | <p>This command adds trace points.</p> <p>The no form of the command removes the trace points.</p> |

router

| | |
|--------------------|--|
| Syntax | router [<i>router-instance</i>] router service-name <i>service-name</i> |
| Context | debug |
| Description | This command enters the context to enable debugging of various protocols and areas of a <i>router-instance</i> . |
| Parameters | <i>router-instance</i> — specifies the router name, CPM router instance, or service ID |
| Values | <i>router-name</i> or <i>service-id</i> |
| | <i>router-instance</i> : <i>router-name</i> <i>router-name</i> Base management <i>cpm-vr-name</i> <i>cpm-vr-name</i> [32 characters maximum] |
| | <i>service-id</i> : 1 to 2147483647 |
| Default | Base |
| | <i>service-name</i> — specifies the service name, up to 64 characters |

ip

| | |
|--------------------|---|
| Syntax | ip |
| Context | debug>router |
| Description | This command configures debugging for IP. |

arp

| | |
|--------------------|--|
| Syntax | arp |
| Context | debug>router>ip |
| Description | This command configures route table debugging. |

icmp

| | |
|----------------|------------------|
| Syntax | [no] icmp |
| Context | debug>router>ip |

Description This command enables ICMP debugging.

icmp6

Syntax **icmp6** [*ip-int-name*]
no icmp6

Context debug>router>ip

Description This command enables ICMPv6 debugging.

interface

Syntax [**no**] **interface** [*ip-int-name* | *ip-address* | ipv6-address | ipv6-address]

Context debug>router>ip

Description This command displays the router IP interface table sorted by interface index.

Parameters *ip-address* — only displays the interface information associated with the specified IP address

Values The following values apply to the 7750 SR and 7950 XRS:

| | |
|--------------|-------------------------------------|
| ipv4-address | a.b.c.d (host bits must be 0) |
| ipv6-address | x:x:x:x:x:x:x (eight 16-bit pieces) |
| | x:x:x:x:x:d.d.d.d |
| | x: [0 to FFFF]H |
| | d: [0 to 255]D |

Values The following values apply to the 7450 ESS:

ipv4-address: a.b.c.d (host bits must be 0)

ip-int-name — only displays the interface information associated with the specified interface name

Values 32 characters maximum

packet

Syntax **packet** [*ip-int-name* | *ip-address*] [**headers**] [*protocol-id*]
no packet [*ip-int-name* | *ip-address*]

Context debug>router>ip

Description This command enables debugging for IP packets.

| | |
|-------------------|--|
| Parameters | <i>ip-int-name</i> — only displays the interface information associated with the specified IP interface name |
| Values | 32 characters maximum |
| | <i>ip-address</i> — only displays the interface information associated with the specified IP address |
| | headers — only displays information associated with the packet header |
| | <i>protocol-id</i> — specifies the decimal value representing the IP protocol to debug. Well-known protocol numbers include ICMP(1), TCP(6), UDP(17). The no form the command removes the protocol from the criteria. |
| Values | 0 to 255 (values can be expressed in decimal, hexadecimal, or binary) |

route-table

| | | | | | | | | | | | | | | | |
|--------------------|--|-------------|-------------------------------|--------------------|---------|-------------|-------------------------------------|--|-------------------|--|-----------------|--|----------------|--------------------|----------|
| Syntax | route-table [<i>ip-prefix/prefix-length</i>] route-table <i>ip-prefix/prefix-length</i> longer no route-table | | | | | | | | | | | | | | |
| Context | debug>router>ip | | | | | | | | | | | | | | |
| Description | This command configures route table debugging. | | | | | | | | | | | | | | |
| Parameters | <i>ip-prefix</i> — the IP prefix for prefix list entry in dotted decimal notation | | | | | | | | | | | | | | |
| Values | The following values apply to the 7750 SR and 7950 XRS: <table><tr><td>ipv4-prefix</td><td>a.b.c.d (host bits must be 0)</td></tr><tr><td>ipv4-prefix-length</td><td>0 to 32</td></tr><tr><td>ipv6-prefix</td><td>x:x:x:x:x:x:x (eight 16-bit pieces)</td></tr><tr><td></td><td>x:x:x:x:x:d.d.d.d</td></tr><tr><td></td><td>x: [0 to FFFF]H</td></tr><tr><td></td><td>d: [0 to 255]D</td></tr><tr><td>ipv6-prefix-length</td><td>0 to 128</td></tr></table> | ipv4-prefix | a.b.c.d (host bits must be 0) | ipv4-prefix-length | 0 to 32 | ipv6-prefix | x:x:x:x:x:x:x (eight 16-bit pieces) | | x:x:x:x:x:d.d.d.d | | x: [0 to FFFF]H | | d: [0 to 255]D | ipv6-prefix-length | 0 to 128 |
| ipv4-prefix | a.b.c.d (host bits must be 0) | | | | | | | | | | | | | | |
| ipv4-prefix-length | 0 to 32 | | | | | | | | | | | | | | |
| ipv6-prefix | x:x:x:x:x:x:x (eight 16-bit pieces) | | | | | | | | | | | | | | |
| | x:x:x:x:x:d.d.d.d | | | | | | | | | | | | | | |
| | x: [0 to FFFF]H | | | | | | | | | | | | | | |
| | d: [0 to 255]D | | | | | | | | | | | | | | |
| ipv6-prefix-length | 0 to 128 | | | | | | | | | | | | | | |
| Values | The following values apply to the 7450 ESS: <table><tr><td>ipv4-prefix</td><td>a.b.c.d (host bits must be 0)</td></tr><tr><td>ipv4-prefix-length</td><td>0 to 32</td></tr></table> | ipv4-prefix | a.b.c.d (host bits must be 0) | ipv4-prefix-length | 0 to 32 | | | | | | | | | | |
| ipv4-prefix | a.b.c.d (host bits must be 0) | | | | | | | | | | | | | | |
| ipv4-prefix-length | 0 to 32 | | | | | | | | | | | | | | |
| | longer — specifies the prefix list entry matches any route that matches the specified <i>ip-prefix</i> and prefix <i>mask</i> length values greater than the specified <i>mask</i> | | | | | | | | | | | | | | |

tunnel-table

| | |
|--------------------|--|
| Syntax | tunnel-table [<i>ip-address</i>] [ldp rsvp [tunnel-id <i>tunnel-id</i>] sdp [sdp-id <i>sdp-id</i>]] |
| Context | debug>router>ip |
| Description | This command enables debugging for tunnel tables. |

l2tp

| | |
|--------------------|---|
| Syntax | l2tp |
| Context | debug>router |
| Description | This command enables the context to configure debugging for L2TP. |

peer

| | |
|--------------------|---|
| Syntax | peer <i>ip-address</i> [{ udp-port <i>port</i> ip }] |
| Context | debug>router>l2tp |
| Description | This command enables and configures debugging for an L2TP peer. |
| Parameters | <i>ip-address</i> — specifies the IP address of the L2TP peer <i>port</i> — specifies the UDP port for the L2TP peer. This parameter is only supported with L2TPv2 peers. ip — displays debugging information for peers using IP transport |

mtrace

| | |
|--------------------|---|
| Syntax | [no] mtrace |
| Context | debug>router |
| Description | This command configures debugging for mtrace. |

misc

| | |
|--------------------|--|
| Syntax | [no] misc |
| Context | debug>router>mtrace |
| Description | This command enables debugging for mtrace miscellaneous. |

packet

| | |
|--------------------|--|
| Syntax | [no] packet [query request response] |
| Context | debug>router>mtrace |
| Description | This command enables debugging for mtrace packets. |

2.13.2.4 Tools Commands

tunnel

| | |
|--------------------|--|
| Syntax | tunnel |
| Context | tools>dump>router>segment-routing> tunnel |
| Description | This command displays Segment Routing tunnels information. |
| Output | |

Sample Output

```
*A:Dut-F#  tools dump router segment-routing tunnel
=====
Legend: (B) - Backup Next-hop for Fast Re-Route
        (D) - Duplicate
=====
-----+
Prefix
|
Sid-Type      Fwd-Type      In-Label  Prot-
Inst
      Next Hop(s)
Label(s) Interface/Tunnel-ID |
-----+
1.0.33.3
Node          Orig/Transit  70000    OSPF-0
              1.0.36.3
              40000          DUTF_TO
_DUTC.1.0      (B)1.0.26.2          30998          DUTF_TO
_DUTB.1.0
1.0.44.4
Node          Orig/Transit  70001    OSPF-0
              1.0.26.2
              30001          DUTF_TO
_DUTB.1.0      (B)1.0.56.5          60001          DUTF_TO
_DUTE.1.0
1.0.55.5
Node          Orig/Transit  70002    OSPF-0
```


| | | | | | |
|-------------|--------------|-------|--------|-------|---------|
| _DUTE.1.0 | 1.0.56.5 | | | 60002 | DUTF_TO |
| _DUTB.1.0 | (B)1.0.26.2 | | | 30995 | DUTF_TO |
| Node | Terminating | 70003 | OSPF-0 | | |
| 1.0.11.1 | | | | | |
| Node | Orig/Transit | 70004 | OSPF-0 | | |
| | 1.0.26.2 | | | 30004 | DUTF_TO |
| _DUTB.1.0 | (B)1.0.36.3 | | | 40004 | DUTF_TO |
| _DUTC.1.0 | | | | | |
| 1.0.22.2 | | | | | |
| Node | Orig/Transit | 70005 | OSPF-0 | | |
| | 1.0.26.2 | | | 30005 | DUTF_TO |
| _DUTB.1.0 | (B)1.0.36.3 | | | 40004 | DUTF_TO |
| _DUTC.1.0 | | | | 20005 | |
| 10.20.1.3 | | | | | |
| Node | Orig/Transit | 70006 | OSPF-0 | | |
| | 1.0.36.3 | | | 40006 | DUTF_TO |
| _DUTC.1.0 | (B)1.0.26.2 | | | 30004 | DUTF_TO |
| _DUTB.1.0 | | | | 20006 | |
| 10.20.1.4 | | | | | |
| Node | Orig/Transit | 70007 | OSPF-0 | | |
| | 1.0.26.2 | | | 30007 | DUTF_TO |
| _DUTB.1.0 | (B)1.0.56.5 | | | 60007 | DUTF_TO |
| _DUTE.1.0 | | | | | |
| 10.20.1.5 | | | | | |
| Node | Orig/Transit | 70008 | OSPF-0 | | |
| | 1.0.56.5 | | | 60008 | DUTF_TO |
| _DUTE.1.0 | (B)1.0.26.2 | | | 30001 | DUTF_TO |
| _DUTB.1.0 | | | | 50008 | |
| Node | Terminating | 70009 | OSPF-0 | | |
| 10.20.1.1 | | | | | |
| Node | Orig/Transit | 70010 | OSPF-0 | | |
| | 1.0.26.2 | | | 30010 | DUTF_TO |
| _DUTB.1.0 | (B)1.0.36.3 | | | 40010 | DUTF_TO |
| _DUTC.1.0 | | | | | |
| 10.20.1.2 | | | | | |
| Node | Orig/Transit | 70011 | OSPF-0 | | |
| | 1.0.26.2 | | | 30011 | DUTF_TO |
| _DUTB.1.0 | (B)1.0.56.5 | | | 60001 | DUTF_TO |
| _DUTE.1.0 | | | | 50011 | |
| Backup Node | Transit | 70994 | OSPF-0 | | |
| | 1.0.56.5 | | | 60994 | DUTF_TO |
| _DUTE.1.0 | | | | | |

```

Backup Node      Transit      70995      OSPF-0
1.0.26.2
30995          DUTF_TO
_DUTB.1.0

Backup Node      Transit      70996      OSPF-0
1.0.26.2
30005          DUTF_TO
_DUTB.1.0

Backup Node      Transit      70998      OSPF-0
1.0.26.2
30998          DUTF_TO
_DUTB.1.0

Backup Node      Transit      70999      OSPF-0
1.0.36.3
40999          DUTF_TO
_DUTC.1.0

Adjacency        Transit      262140     OSPF-0
1.0.26.2
3              DUTF_TO
_DUTB.1.0
(B)1.0.36.3
40004          DUTF_TO
_DUTC.1.0
20005

Adjacency        Transit      262141     OSPF-0
1.0.56.5
3              DUTF_TO
_DUTE.1.0

Adjacency        Transit      262142     OSPF-0
1.0.36.3
3              DUTF_TO
_DUTC.1.0

Adjacency        Transit      262143     OSPF-0
1.0.26.2
3              DUTF_TO
_DUTB.1.0
(B)1.0.56.5
60001          DUTF_TO
_DUTE.1.0
50011

*A:Dut-F#

*A:Dut-A# tools dump router segment-routing tunnel
=====
Legend: (B) - Backup Next-hop for Fast Re-
Route
(D) Duplicate
=====
-----
Prefix
|
Sid-Type      Fwd-Type      In-Label  Prot-
Inst
      Next Hop(s)
Label(s) Interface/Tunnel-ID | Out-
-----

Adjacency      Transit      262136     ISIS-0
10.10.2.1      10.10.2.3
3              ip-

```

```

Adjacency      Transit      262137   ISIS-0
10.10.2.1      10.10.2.3                      3      ip-

Adjacency      Transit      262138   ISIS-0
10.10.1.1      10.10.1.2                      3      ip-

Adjacency      Transit      262139   ISIS-0
10.10.1.1      10.10.1.2                      3      ip-

Node           Terminating 474387   ISIS-0
10.20.1.2
Node           Orig/Transit 474388   ISIS-0
10.10.1.1      10.10.1.2                      474388  ip-
10.20.1.3
Node           Orig/Transit 474389   ISIS-0
10.10.2.1      10.10.2.3                      474389  ip-
10.20.1.4
Node           Orig/Transit 475287   ISIS-0
10.10.1.1      10.10.1.2                      475287  ip-
10.20.1.5
Node           Orig/Transit 475288   ISIS-0
10.10.2.1      10.10.2.3                      475288  ip-
10.20.1.6
Node           Orig/Transit 475289   ISIS-0
10.10.1.1      10.10.1.2                      475289  ip-
*A:Dut-A#

```

*A:Dut-C# tools dump router segment-routing tunnel

Legend: (B) - Backup Next-hop for Fast Re-Route

(D) - Duplicate

Prefix

| Sid-Type | Fwd-Type | In-Label | Prot- | Out- |
|----------|---------------------|----------|-------|------|
| Inst | Next Hop(s) | | | |
| Label(s) | Interface/Tunnel-ID | | | |

```

Adjacency      Transit      262129   ISIS-0
10.10.12.3     10.10.12.2                      3      ip-
(B)10.10.3.2   10.10.3.2                      3      ip-

Adjacency      Transit      262130   ISIS-0
10.10.12.3     10.10.12.2                      3      ip-

```

| | | | | | |
|------------|----------------|--------|--------|--------|-----|
| 10.10.12.3 | (B) 10.10.3.2 | | | 3 | ip- |
| 10.10.3.3 | | | | | |
| Adjacency | Transit | 262133 | ISIS-0 | | |
| | 10.10.5.5 | | | 3 | ip- |
| 10.10.5.3 | (B) 10.10.12.2 | | | 474389 | ip- |
| 10.10.12.3 | | | | 474390 | |
| Adjacency | Transit | 262134 | ISIS-0 | | |
| | 10.10.5.5 | | | 3 | ip- |
| 10.10.5.3 | (B) 10.10.12.2 | | | 474389 | ip- |
| 10.10.12.3 | | | | 474390 | |
| Adjacency | Transit | 262135 | ISIS-0 | | |
| | 10.10.3.2 | | | 3 | ip- |
| 10.10.3.3 | (B) 10.10.12.2 | | | 3 | ip- |
| 10.10.12.3 | | | | | |
| Adjacency | Transit | 262136 | ISIS-0 | | |
| | 10.10.3.2 | | | 3 | ip- |
| 10.10.3.3 | (B) 10.10.12.2 | | | 3 | ip- |
| 10.10.12.3 | | | | | |
| Adjacency | Transit | 262137 | ISIS-0 | | |
| | 10.10.2.1 | | | 3 | ip- |
| 10.10.2.3 | | | | | |
| Adjacency | Transit | 262138 | ISIS-0 | | |
| | 10.10.2.1 | | | 3 | ip- |
| 10.10.2.3 | | | | | |
| 10.20.1.4 | Orig/Transit | 474389 | ISIS-0 | | |
| Node | 10.10.12.2 | | | 474389 | ip- |
| 10.10.12.3 | (B) 10.10.5.5 | | | 474389 | ip- |
| 10.10.5.3 | | | | | |
| 10.20.1.5 | Orig/Transit | 474390 | ISIS-0 | | |
| Node | 10.10.5.5 | | | 474390 | ip- |
| 10.10.5.3 | (B) 10.10.12.2 | | | 474389 | ip- |
| 10.10.12.3 | | | | 474390 | |
| 10.20.1.6 | Orig/Transit | 474391 | ISIS-0 | | |
| Node | 10.10.5.5 | | | 474391 | ip- |
| 10.10.5.3 | (B) 10.10.12.2 | | | 474391 | ip- |
| 10.10.12.3 | | | | | |
| 10.20.1.2 | Orig/Transit | 474392 | ISIS-0 | | |
| Node | 10.10.12.2 | | | 474392 | ip- |

```

10.10.12.3
(B)10.10.3.2 474392 ip-
10.10.3.3

Node Terminating 474393 ISIS-0
*A:Dut-C#

*A:Dut-C# tools dump router segment-routing tunnel
=====
Legend: (B) - Backup Next-hop for Fast Re-
Route
(D) Duplicate
=====
-----
Prefix
|
Sid-Type Fwd-Type In-Label Prot-
Inst Next Hop(s) Out-
Label(s) Interface/Tunnel-ID |
-----

Adjacency Transit 262129 ISIS-0
10.10.12.3 10.10.12.2 3 ip-
(B)10.10.3.2 3 ip-
10.10.3.3

Adjacency Transit 262130 ISIS-0
10.10.12.3 10.10.12.2 3 ip-
(B)10.10.3.2 3 ip-
10.10.3.3

Adjacency Transit 262133 ISIS-0
10.10.5.3 10.10.5.5 3 ip-
(B)10.10.12.2 474389 ip-
10.10.12.3 474390

Adjacency Transit 262134 ISIS-0
10.10.5.3 10.10.5.5 3 ip-
(B)10.10.12.2 474389 ip-
10.10.12.3 474390

Adjacency Transit 262135 ISIS-0
10.10.3.3 10.10.3.2 3 ip-
(B)10.10.12.2 3 ip-
10.10.12.3

Adjacency Transit 262136 ISIS-0
10.10.3.3 10.10.3.2 3 ip-
(B)10.10.12.2 3 ip-
10.10.12.3

Adjacency Transit 262137 ISIS-0

```

```

10.10.2.3          10.10.2.1          3          ip-
Adjacency          Transit          262138    ISIS-0
10.10.2.3          10.10.2.1          3          ip-
10.20.1.4
Node              Orig/Transit      474389    ISIS-0
10.10.12.3         10.10.12.2         474389    ip-
(B)10.10.5.5       474389    ip-
10.10.5.3
10.20.1.5
Node              Orig/Transit      474390    ISIS-0
10.10.5.3         10.10.5.5         474390    ip-
(B)10.10.12.2     474389    ip-
10.10.12.3         474390
10.20.1.6
Node              Orig/Transit      474391    ISIS-0
10.10.5.3         10.10.5.5         474391    ip-
(B)10.10.12.2     474391    ip-
10.10.12.3
10.20.1.2
Node              Orig/Transit      474392    ISIS-0
10.10.12.3         10.10.12.2         474392    ip-
(B)10.10.3.2     474392    ip-
10.10.3.3
Node              Terminating      474393    ISIS-0
*A:Dut-C#

```

l2tp

Syntax **l2tp**

Context tools>perform>router

Description This command enables the context to configure performance tools for L2TP.

peer

Syntax **peer** *ip-address* [{**udp-port** *port* | **ip**}]

Context tools>perform>router>l2tp

Description This command configures performance tools for an L2TP peer.

Parameters *ip-address* — specifies the IP address of the L2TP peer.

port — specifies the UDP port for the L2TP peer. This parameter is only supported with L2TPv2 peers.

ip — enables performance tools for peers using IP transport.

3 VRRP

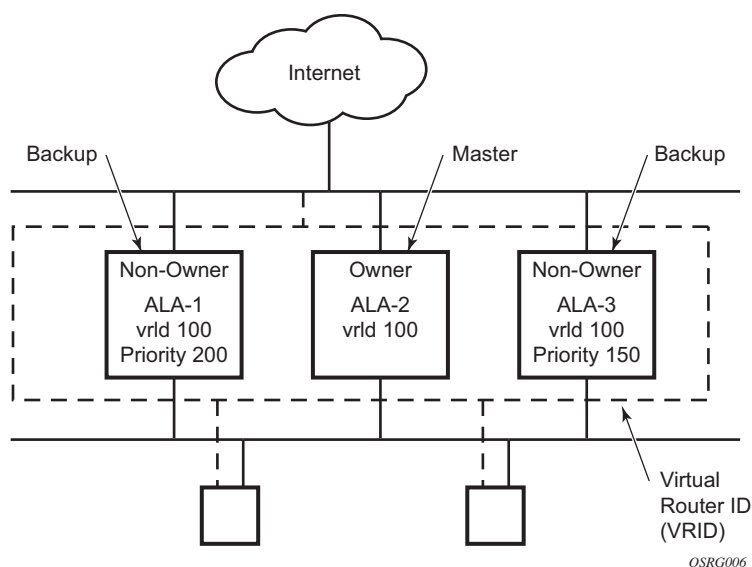
3.1 VRRP Overview

The Virtual Router Redundancy Protocol (VRRP) for IPv4 is defined in the IETF RFC 3768, *Virtual Router Redundancy Protocol*. VRRP for IPv6 is specified in *draft-ietf-vrrp-unified-spec-02.txt* and only applies to the 7750 SR and 7950 XRS. VRRP describes a method of implementing a redundant IP interface shared between two or more routers on a common LAN segment, allowing a group of routers to function as one virtual router. When this IP interface is specified as a default gateway on hosts directly attached to this LAN, the routers sharing the IP interface prevent a single point of failure by limiting access to this gateway address. VRRP can be implemented on IES service interfaces and on core network IP interfaces.

If the master virtual router fails, the backup router configured with the highest acceptable priority becomes the master virtual router. The new master router assumes the normal packet forwarding for the local hosts.

Figure 23 shows an example of a VRRP configuration.

Figure 23 VRRP Configuration



3.2 VRRP Components

VRRP consists of the following components:

3.2.1 Virtual Router

A virtual router is a logical entity managed by VRRP that acts as a default router for hosts on a shared LAN. It consists of a Virtual Router Identifier (VRID) and a set of associated IP addresses (or an address) across a common LAN. A VRRP router can be the backup for one or more virtual routers.

The purpose of supporting multiple IP addresses within a single virtual router is for multi-netting. This is a common mechanism that allows multiple local subnet attachments on a single routing interface. Up to four virtual routers are possible on a single Nokia IP interface. The virtual routers must be in the same subnet. Each virtual router has its own VRID, state machine, and messaging instance.

3.2.2 IP Address Owner

VRRP can be configured in either an owner or non-owner mode. The owner is the VRRP router whose virtual router IP address is the same as the real interface IP address. This is the router that responds to packets addressed to one of the IP addresses for ICMP pings, TCP connections, and so on. All other virtual router instances participating in this message domain must have the same VRID configured and cannot be configured as owner.

Nokia routers allow the virtual routers to be configured as non-owners of the IP address. VRRP on a router can be configured to allow non-owners to respond to ICMP echo requests when they become the virtual router master for the virtual router. Telnet and other connection-oriented protocols can also be configured for non-owner master response. However, the individual application conversations (connections) will not survive a VRRP failover. A non-owner VRRP router operating as a backup will not respond to any packets addressed to any of the virtual router IP addresses.

3.2.3 Primary and Secondary IP Addresses

A primary address is an IP address selected from the set of real interface address. VRRP advertisements are always sent using the primary IP address as the source of the IP packet.

An IP interface must always have a primary IP address assigned for VRRP to be active on the interface. Nokia routers supports both primary and secondary IP addresses (multi-netting) on the IP interface. The virtual router's VRID primary IP address is always the primary address on the IP interface. VRRP uses the primary IP address as the IP address placed in the source IP address field of the IP header for all VRRP messages sent on that interface.

3.2.4 Virtual Router Master

The VRRP router that controls the IP addresses associated with a virtual router is called the master. The master is responsible for forwarding packets sent to the VRRP IP addresses. An election process provides dynamic failover of the forwarding responsibility if the master becomes unavailable. This allows any of the virtual router IP addresses on the LAN to be used as the default first hop router by end hosts. This enables a higher availability default path without requiring configuration of dynamic routing or router discovery protocols on every end host.

If the master is unavailable, each backup virtual router for the VRID compares the configured priority values to determine the master role. In case of a tie, the virtual router with the highest primary IP address becomes master.

The preempt parameter can be set to false to prevent a backup virtual router with a better priority value from becoming master when an existing non-owner virtual router is the current master. This is determined on a first-come, first-served basis.

While master, a virtual router routes and originates all IP packets into the LAN using the physical MAC address for the IP interface as the Layer 2 source MAC address, not the VRID MAC address. ARP packets also use the parent IP interface MAC address as the Layer 2 source MAC address while inserting the virtual router MAC address in the appropriate hardware address field. VRRP messages are the only packets transmitted using the virtual router MAC address as the Layer 2 source MAC address.

3.2.5 Virtual Router Backup

A new virtual router master is selected from the set of VRRP routers available to assume forwarding responsibility for a virtual router in case the current master fails.

3.2.6 Owner and Non-Owner VRRP

The owner controls the IP address of the virtual router and is responsible for forwarding packets sent to this IP address. The owner assumes the role of the master virtual router. Only one virtual router in the domain can be configured as owner. All other virtual router instances participating in this message domain must have the same VRID configured.

The most important parameter to be defined on a non-owner virtual router instance is the priority. The priority defines a virtual router's selection order in the master election process. The priority value and the preempt mode determine the virtual router with the highest priority to become the master virtual router.

The base priority is used to determine the in-use priority of the virtual router instance as modified by any optional VRRP priority control policy. VRRP priority control policies can be used to either override or adjust the base priority value depending on events or conditions within the chassis.

For information about non-owner access parameters, see [VRRP Non-Owner Accessibility](#).

3.2.7 Configurable Parameters

As well as to backup IP addresses, to facilitate configuration of a virtual router on Nokia routers, the following parameters can be defined in owner configurations:

- [Virtual Router ID \(VRID\)](#)
- [Message Interval and Master Inheritance](#)
- [VRRP Message Authentication](#)
- [Authentication Data](#)
- [Virtual MAC Address](#)

The following parameters can be defined in non-owner configurations:

- [Virtual Router ID \(VRID\)](#)

- [Priority](#)
- [Message Interval and Master Inheritance](#)
- [Master Down Interval](#)
- [Preempt Mode](#)
- [VRRP Message Authentication](#)
- [Authentication Data](#)
- [Virtual MAC Address](#)
- [Inherit Master VRRP Router's Advertisement Interval Timer](#)
- [Policies](#)

3.2.7.1 Virtual Router ID (VRID)

The VRID must be configured with the same value on each virtual router associated with the redundant IP address (or addresses). The VRID is placed in all VRRP advertisement messages sent by each virtual router.

3.2.7.2 Priority

The priority value affects the interaction between this VRID and the same VRID of other virtual routers participating on the same LAN. A higher-priority value defines a greater priority in becoming the virtual router master for the VRID. The priority value can only be configured when the defined IP address on the IP interface is different from the virtual router IP address (non-owner mode).

When the IP address on the IP interface matches the virtual router IP address (owner mode), the priority value is fixed at 255, the highest value possible. This virtual router member is considered the owner of the virtual router IP address. There can only be one owner of the virtual router IP address for all virtual router members.

The priority value 0 is reserved for VRRP advertisement message purposes. It is used to tell other virtual routers in the same VRID that this virtual router is no longer acting as master, triggering a new election process. When this happens, each backup virtual router sets its master down timer equal to the skew time value. This shortens the time until one of the backup virtual routers becomes master.

The current master virtual router must transmit a VRRP advertisement message immediately upon receipt of a VRRP message with priority set to 0. This prevents another backup from becoming master for a short period of time.

Non-owner virtual routers may be configured with a priority of 254 through 1. The default value is 100. Multiple non-owners can share the same priority value. When multiple non-owner backup virtual routers are tied (transmit VRRP advertisement messages simultaneously) in the election process, all attempt to become master simultaneously; the one with the best priority will win the election. If the priority value in the message is equal to the master's local priority value, the primary IP address of the local master and of the message is evaluated as the tie breaker. The higher IP address becomes master. (The primary IP address is the source IP address of the VRRP advertisement message.)

The priority value is also used to determine when to preempt the existing master. If the preempt mode value is true, VRRP advertisement messages from inferior (lower-priority) masters are discarded, causing the master down timer to expire and causing the transition to master state.

The priority value also dictates the skew time added to the master timeout period.

3.2.7.3 IP Addresses

Each virtual router with the same VRID should be defined with the same set of IP addresses. These are the IP addresses being used by hosts on the LAN as gateway addresses. Multi-netting supports 16 IP addresses on the IP interface; up to 16 addresses can be assigned to a specific virtual router instance.

3.2.7.4 Message Interval and Master Inheritance

Each virtual router is configured with a message interval per VRID within which it participates. This parameter must be the same for every virtual router on the VRID.

For IPv4, the default advertisement interval is 1 s and can be configured between 100 ms and 255 s 900 ms. For IPv6, the default advertisement interval is 1 s and can be configured between 100 ms and 40 s 950 ms.

As specified in the RFC, the advertisement interval field in every received VRRP advertisement message must match the locally configured advertisement interval. If a mismatch occurs, depending on the inherit configuration, the current master's advertisement interval setting can be used to operationally override the locally configured advertisement interval setting. If the current master changes, the new master setting is used. If the local virtual router becomes master, the locally configured advertisement interval is enforced.

If a VRRP advertisement message is received with an advertisement interval set to a value different from the local value and the inherit parameter is disabled, the message is discarded without processing.

The master virtual router on a VRID uses the advertisement interval to load the advertisement timer, specifying when to send the next VRRP advertisement message. Each backup virtual router on a VRID uses the advertisement interval (with the configured local priority) to determine the master down timer value.

VRRP advertisement messages that are fragmented, or contain IP options (IPv4), or contain extension headers (IPv6) require a longer message interval to be configured.

3.2.7.5 Skew Time

The skew time is used to add a time period to the master down interval. This is not a configurable parameter. It is determined from the current local priority of the virtual router's VRID. To calculate the skew time, the virtual router evaluates the following formula:

For IPv4: $\text{Skew Time} = ((256 - \text{priority}) / 256) \text{ seconds}$

For IPv6: $\text{Skew Time} = (((256 - \text{priority}) * \text{Master_Adver_Interval}) / 256) \text{ centiseconds}$

The higher the priority value, the shorter the skew time will be. This means that virtual routers with a lower priority will transition to master slower than virtual routers with a higher priority.

3.2.7.6 Master Down Interval

The master down interval is a calculated value used to load the master down timer. When the master down timer expires, the virtual router enters the master state. To calculate the master down interval, the virtual router evaluates the following formula:

$\text{Master Down Interval} = (3 \times \text{Operational Advertisement Interval}) + \text{Skew Time}$

The operational advertisement interval is dependent upon the state of the inherit parameter. When the inherit parameter is enabled, the operational advertisement interval is determined from the current master's advertisement interval field in the VRRP advertisement message. When inherit is disabled, the operational advertisement interval must be equal to the locally configured advertisement interval.

The master down timer is only operational when the local virtual router is operating in backup mode.

3.2.7.7 Preempt Mode

Preempt mode is a true or false configured value that controls whether a specific backup virtual router preempts a lower-priority master. The IP address owner will always become master when available. Preempt mode cannot be set to false on the owner virtual router. The default value for preempt mode is true.

When preempt mode is true, a master non-owner virtual router will only allow itself to be preempted when the incoming VRRP advertisement message priority field value is one of the following:

- Greater than the virtual router in-use priority value
- Equal to the in-use priority value, and the source IP address (primary IP address) is greater than the virtual router instance primary IP address

A backup router will only attempt to become the master router if the preempt mode is true and the received VRRP advertisement priority field is less than the virtual router in-use priority value.

3.2.7.8 VRRP Message Authentication

The authentication type parameter defines the type of authentication used by the virtual router in VRRP advertisement message authentication. VRRP message authentication is applicable to IPv4 only. The current master uses the configured authentication type to indicate any egress message manipulation that must be performed in conjunction with any supporting authentication parameters before transmitting a VRRP advertisement message. The configured authentication type value is transmitted in the message authentication type field with the appropriate authentication data field filled in. Backup routers use the authentication type message field value in interpreting the contained authentication data field within received VRRP advertisement messages.

VRRP supports three message authentication methods that provide varying degrees of security. The supported authentication types are:

- 0 – No Authentication
- 1 – Simple Text Password
- 2 – IP Authentication Header

3.2.7.8.1 Authentication Type 0 – No Authentication

The use of authentication type 0 indicates that VRRP advertisement messages are not authenticated (provides no authentication). The master transmitting VRRP advertisement messages will transmit the value 0 in the egress messages authentication type field and the authentication data field. Backup virtual routers receiving VRRP advertisement messages with the authentication type field equal to 0 will ignore the authentication data field in the message.

All compliant VRRP advertisement messages are accepted. The following fields within the received VRRP advertisement message are checked for compliance (the VRRP specification may require additional checks):

- IP header checks specific to VRRP
 - IP header destination IP address – Must be 224.0.0.18
 - IP header TTL field – Must be equal to 255; the packet must not have traversed any IP routed hops
 - IP header protocol field – must be 112 (decimal)
- VRRP message checks
 - Version field – Must be set to the value of 2
 - Type field – Must be set to the value of 1 (advertisement)
 - Virtual router ID field – Must match one of the configured VRIDs on the ingress IP interface (all other fields are dependent on matching the virtual router ID field to one of the interfaces configured *VRID* parameters)
 - Priority field – Must be equal to or greater than the *VRID* in-use priority or be equal to 0 (if equal to the *VRID* in-use priority and 0, requires further processing regarding master/backup and senders IP address to determine validity of the message)
 - Authentication type field – Must be equal to 0
 - Advertisement interval field – Must be equal to the *VRID* configured advertisement interval
 - Checksum field – Must be valid
 - Authentication data fields – Must be ignored

VRRP messages not meeting the criteria are silently discarded.

3.2.7.8.2 Authentication Type 1 – Simple Text Password

The use of authentication type 1 indicates that VRRP advertisement messages are authenticated with a clear (simple) text password. All virtual routers participating in the virtual router instance must be configured with the same 8 octet password. Transmitting virtual routers put a value of 1 in the VRRP advertisement message authentication type field and put the configured simple text password into the message authentication data field. Receiving virtual routers compare the message authentication data field with the local configured simple text password based on the message authentication type field value of 1.

The same checks are performed as for type 0, with the following exceptions (the VRRP specification may require additional checks):

- VRRP message checks
 - Authentication type field – Must be equal to 1
 - Authentication data fields – Must be equal to the *VRID* configured simple text password

Any VRRP messages not meeting the type 0 verification checks, with the preceding exceptions are silently discarded.

3.2.7.8.3 Authentication Failure

Any received VRRP advertisement message that fails authentication must be silently discarded with an invalid authentication counter incremented for the ingress virtual router instance.

3.2.7.9 Authentication Data

This feature is different from the VRRP advertisement message field with the same name. Authentication data is any required authentication information that is pertinent to the configured authentication type. The type of authentication data used for each authentication type is listed in [Table 34](#).

Table 34 Authentication Data Type

| Authentication Type | Authentication Data |
|---------------------|---------------------------------------|
| 0 | None, authentication is not performed |

Table 34 Authentication Data Type (Continued)

| Authentication Type | Authentication Data |
|---------------------|---|
| 1 | Simple text password consisting of 8 octets |

3.2.7.10 Virtual MAC Address

The MAC address can be used instead of an IP address in ARP responses when the virtual router instance is master. The MAC address configuration must be the same for all virtual routers participating as a virtual router, or indeterminate connectivity by the attached IP hosts will result. All VRRP advertisement messages are transmitted with *ieee-mac-addr* as the source MAC.

3.2.7.11 VRRP Advertisement Message IP Address List Verification

VRRP advertisement messages contain an IP address count field that indicates the number of IP addresses listed in the sequential IP address fields at the end of the message.

The Nokia routers implementation always logs mismatching events. The decision on where and whether to forward the generated messages depends on the configuration of the event manager.

To facilitate the sending of mismatch log messages, each virtual router instance keeps the mismatch state associated with each source IP address in the VRRP master table. Whenever the state changes, a mismatch log message is generated indicating the source IP address within the message, the mismatch or match event, and the time of the event.

With secondary IP address support, multiple IP addresses can be in the list and each should match the IP address on the virtual router instance. Owner and non-owner virtual router instances have the supported IP addresses explicitly defined, making mismatched supported IP addresses within the interconnected virtual router instances a provisioning issue.

3.2.7.12 Inherit Master VRRP Router's Advertisement Interval Timer

The virtual router instance can inherit the master VRRP router's advertisement interval timer, which is used by backup routers to calculate the master down timer.

The inheritance is only configurable in the non-owner nodal context. The inheritance is used to allow the current virtual router instance master to dictate the master down timer for all backup virtual routers.

3.2.7.13 IPv6 Virtual Router Instance Operationally Up

After the 7750 SR or 7950 XRS IPv6 virtual router is configured with a minimum of one link-local backup address, the parent interface's router advertisement must be configured to use the virtual MAC address for the virtual router to be considered operationally up.

3.2.7.14 Policies

Policies can be configured to control VRRP priority with the virtual router instance. VRRP priority control policies can be used to override or adjust the base priority value, depending on events or conditions within the chassis.

The policy can be associated with more than one virtual router instance. The priority events within the policy override or diminish the base priority dynamically affecting the in-use priority. As priority events clear in the policy, the in-use priority can eventually be restored to the base priority value.

Policies can only be configured in the non-owner VRRP context. For non-owner virtual router instances, if policies are not configured, then the base priority is used as the in-use priority.

3.3 VRRP Priority Control Policies

This implementation of VRRP supports control policies to manipulate virtual router participation in the VRRP master election process and master self-deprecation. The local priority value for the virtual router instance is used to control the election process and master state.

3.3.1 VRRP Virtual Router Policy Constraints

Priority control policies can only be applied to non-owner VRRP virtual router instances. Owner VRRP virtual routers cannot be controlled by a priority control policy because they are required to have a priority value of 255 that cannot be diminished. Only one VRRP priority control policy can be applied to a non-owner virtual router instance.

Multiple VRRP virtual router instances may be associated with the same IP interface, allowing multiple priority control policies to be associated with the IP interface.

An applied VRRP priority control policy only affects the in-use priority on the virtual router instance when the preempt mode has been enabled. A virtual router instance with preempt mode disabled will always use the base priority as the in-use priority, ignoring any configured priority control policy.

3.3.2 VRRP Virtual Router Instance Base Priority

Non-owner virtual router instances must have a base priority value between 1 and 254. The value 0 is reserved for master termination. The value 255 is reserved for owners. The default base priority for non-owner virtual router instances is the value 100.

The base priority is the starting priority for the VRRP instance. The actual in-use priority for the VRRP instance is determined from the base priority and an optional VRRP priority control policy.

3.3.3 VRRP Priority Control Policy Delta In-Use Priority Limit

A VRRP priority control policy enforces an overall minimum value that the policy can assign to the VRRP virtual router instance base priority. This value provides a lower limit to the delta priority events manipulation of the base priority.

A delta priority event is a conditional event defined in the priority control policy that subtracts a specified amount from the current, in-use priority for all VRRP virtual router instances to which the policy is applied. Multiple delta priority events can apply simultaneously, creating a dynamic priority value. The base priority for the instance, less the sum of the delta values, determines the priority value in-use.

An explicit priority event is a conditional event defined in the priority control policy that explicitly defines the in-use priority value for the virtual router instance. The explicitly defined value is not affected by the delta in-use priority limit. When multiple explicit priority events happen simultaneously, the lowest value is used for the in-use priority. The configured base priority is not a factor in explicit priority overrides of the in-use priority.

The allowed range of the Delta In-Use Priority Limit is 1 to 254. The default is 1, which prevents the delta priority events from operationally disabling the virtual router instance.

3.3.4 VRRP Priority Control Policy Priority Events

The main function of a VRRP priority control policy is to define conditions or events that affect the system's ability to communicate with outside hosts or portions of the network. When one or multiple of these events are true, the base priority on the virtual router instance is either overwritten with an explicit value, or a sum of delta priorities is subtracted from the base priority. The result is the in-use priority for the virtual router instance. Any priority event may be configured as an explicit event or a delta event.

Explicit events override all delta events. When multiple explicit events occur, the event with the lowest priority value is assigned to the in-use priority. As events clear, the in-use priority is reevaluated accordingly and adjusted dynamically.

Delta priority events also have priority values. When no explicit events have occurred within the policy, the sum of the occurring delta events priorities is subtracted from the base priority of each virtual router instance. If the result is lower than the delta in-use priority limit, the delta in-use priority limit is used as the in-use priority for the virtual router instance. Otherwise, the in-use priority is set to the base priority less the sum of the delta events.

Each event generates a VRRP priority event message indicating the policy-id, the event type, the priority type (delta or explicit), and the event priority value. Another log message is generated when the event is no longer true, indicating that it has been cleared.

3.3.4.1 Priority Event Hold-Set Timers

Hold-set timers are used to dampen the effect of a flapping event. A flapping event is where the event continually transitions between clear and set. The hold-set value is loaded into a hold-set timer that prevents a set event from transitioning to the cleared state until it expires.

Each time an event transitions between cleared and set, the timer is loaded and begins to count down to zero. If the timer reaches zero, the event will be allowed to enter the cleared state again. Entering the cleared state is always dependent on the object controlling the event conforming to the requirements defined in the event. It is possible, on some event types, to have a further set action reload the hold-set timer. This extends the time that must pass before the hold-set timer expires, and the event enters the cleared state.

For an example of a hold-set timer setting, refer to [LAG Degrade Priority Event](#).

3.3.4.2 Port Down Priority Event

The port down priority event is assigned to either a physical port or a SONET/SDH channel for the 7750 SR and 7450 ESS. The port or channel operational state is evaluated to determine a port down priority event or event clear.

When the port or channel operational state is up, the port down priority event is considered false or cleared. When the port or channel operational state is down, the port down priority event is considered true or set.

3.3.4.3 LAG Degrade Priority Event

The LAG degrade priority event is assigned to an existing Link Aggregation Group (LAG). The LAG degrade priority event is conditional on a percentage of available port bandwidth on the LAG. Multiple bandwidth percentage thresholds may be defined, each with its own priority value.

If the LAG transitions from one threshold to the next, the previous threshold priority value is subtracted from the total delta sum while the new threshold priority value is added to the sum. The new sum is then subtracted from the base priority and compared to the delta in-use priority limit to determine the new in-use priority on the virtual router instance.

The following example shows a LAG degrade priority event and its interaction with the hold-set timer in changing the in-use priority.

The following state and timer settings are used for the LAG events listed in [Table 35](#):

- User-defined thresholds: 2 ports down, 4 ports down, 6 ports down
- LAG configured ports: 8 ports
- Hold-set timer (hold-set): 5 seconds

Table 35 LAG Events

| Time (seconds) | LAG Port State | Parameter | State | Comments |
|----------------|----------------|-----------------|--------------------|--|
| 0 | All ports down | Event State | Set - 8 ports down | — |
| | | Event Threshold | 6 ports down | — |
| | | Hold-set Timer | 5 seconds | Set to hold-set parameter |
| 1 | One port up | Event State | Set - 8 ports down | Cannot change until hold-set timer expires |
| | | Event Threshold | 6 ports down | — |
| | | Hold-set Timer | 5 seconds | Event does not affect timer |
| 2 | All ports up | Event State | Set - 8 ports down | Still waiting for hold-set timer expiry |
| | | Event Threshold | 6 ports down | — |
| | | Hold-set Timer | 3 seconds | — |

Table 35 LAG Events (Continued)

| Time (seconds) | LAG Port State | Parameter | State | Comments |
|----------------|------------------|-----------------|------------------------|--|
| 5 | All ports up | Event State | Cleared - All ports up | — |
| | | Event Threshold | None | Event cleared |
| | | Hold-set Timer | Expired | — |
| 100 | Five ports down | Event State | Set - 5 ports down | — |
| | | Event Threshold | 4 ports down | — |
| | | Hold-set Timer | Expired | Set to hold-set parameter |
| 102 | Three ports down | Event State | Set - 5 ports down | — |
| | | Event Threshold | 4 ports down | — |
| | | Hold-set Timer | 3 seconds | — |
| 103 | All ports up | Event State | Set - 5 ports down | — |
| | | Event Threshold | 4 ports down | — |
| | | Hold-set Timer | 2 second | — |
| 104 | Two ports down | Event State | Set - 5 ports down | — |
| | | Event Threshold | 4 ports down | — |
| | | Hold-set timer | 1 second | Current threshold is 5, so 2 down has no effect |
| 105 | Two ports down | Event State | Set - 2 ports down | — |
| | | Event Threshold | 2 ports down | — |
| | | Hold-set timer | Expired | — |
| 200 | Four ports down | Event State | Set - 2 ports down | — |
| | | Event Threshold | 4 ports down | — |
| | | Hold-set timer | 5 seconds | Set to hold-set parameter |
| 202 | Seven ports down | Event State | Set - 7 ports down | Changed due to increase |
| | | Event Threshold | 6 ports down | — |
| | | Hold-set timer | 5 seconds | Set to hold-set due to threshold increase |

Table 35 LAG Events (Continued)

| Time (seconds) | LAG Port State | Parameter | State | Comments |
|----------------|----------------|-----------------|------------------------|---------------|
| 206 | All ports up | Event State | Set - 7 ports down | — |
| | | Event Threshold | 6 ports down | — |
| | | Hold-set timer | 1 second | — |
| 207 | All ports up | Event State | Cleared - All ports up | — |
| | | Event Threshold | None | Event cleared |
| | | Hold-set timer | Expired | — |

3.3.4.4 Host Unreachable Priority Event

The host unreachable priority event creates a continuous ping task that is used to test connectivity to a remote host. The path to the remote host and the remote host itself must be capable and configured to accept ICMP echo request and replies for the ping to be successful.

The ping task is controlled by interval and size parameters that define how often the ICMP request messages are transmitted and the size of each message. A historical missing reply parameter defines when the ping destination is considered unreachable.

When the host is unreachable, the host unreachable priority event is considered true or set. When the host is reachable, the host unreachable priority event is considered false or cleared.

3.3.4.5 Route Unknown Priority Event

The route unknown priority event defines a task that monitors the existence of a specific route prefix in the system's routing table.

The route monitoring task can be constrained by a condition that allows a prefix, which is less specific than the defined prefix, to be considered as a match. The source protocol can be defined to indicate the protocol that the installed route must be populated from. To further define match criteria when multiple instances of the route prefix exist, an optional next hop parameter can be defined.

When a route prefix exists within the active route table that matches the defined match criteria, the route unknown priority event is considered false or cleared. When a route prefix does not exist within the active route table matching the defined criteria, the route unknown priority event is considered true or set.

3.4 VRRP Non-Owner Accessibility

Although the RFC states that only VRRP owners can respond to ping and other management-oriented protocols directed to the VRID IP addresses, the routers allow an override of this restraint on a per VRRP virtual router instance basis.

3.4.1 Non-Owner Access Ping Reply

When non-owner access ping reply is enabled on a virtual router instance, ICMP echo request messages destined to the non-owner virtual router instance IP addresses are not discarded at the IP interface when operating in master mode. ICMP echo request messages are always discarded in backup mode.

When non-owner access ping reply is disabled on a virtual router instance, ICMP echo request messages destined to the non-owner virtual router instance IP addresses are silently discarded in both the master and backup modes.

3.4.2 Non-Owner Access Telnet

When non-owner access Telnet is enabled on a virtual router instance, authorized Telnet sessions may be established that are destined to the virtual router instance IP addresses when operating in master mode. Telnet sessions are always discarded at the IP interface when destined to a virtual router IP address operating in backup mode. Enabling non-owner access Telnet does not guarantee Telnet access; correct management and security features must be enabled to allow Telnet on this interface and possibly from the specified source IP address.

When non-owner access Telnet is disabled on a virtual router instance, Telnet sessions destined to the non-owner virtual router instance IP addresses are silently discarded in both master and backup modes.

3.4.3 Non-Owner Access SSH

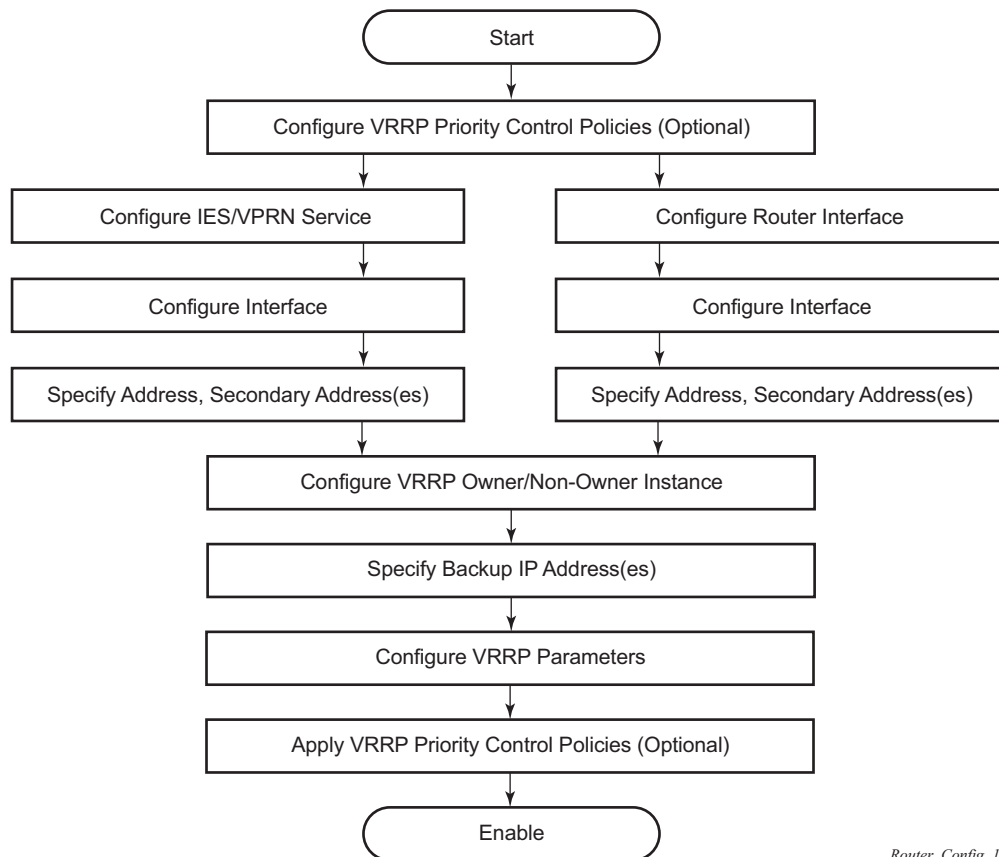
When non-owner access SSH is enabled on a virtual router instance, authorized SSH sessions may be established that are destined to the virtual router instance IP addresses when operating in master mode. SSH sessions are always discarded at the IP interface when destined to a virtual router IP address operating in backup mode. Enabling non-owner access SSH does not guarantee SSH access; correct management and security features must be enabled to allow SSH on this interface and possibly from the specified source IP address. SSH is applicable to IPv4 VRRP only.

When non-owner access SSH is disabled on a virtual router instance, SSH sessions destined to the non-owner virtual router instance IP addresses are silently discarded in both master and backup modes.

3.5 VRRP Configuration Process Overview

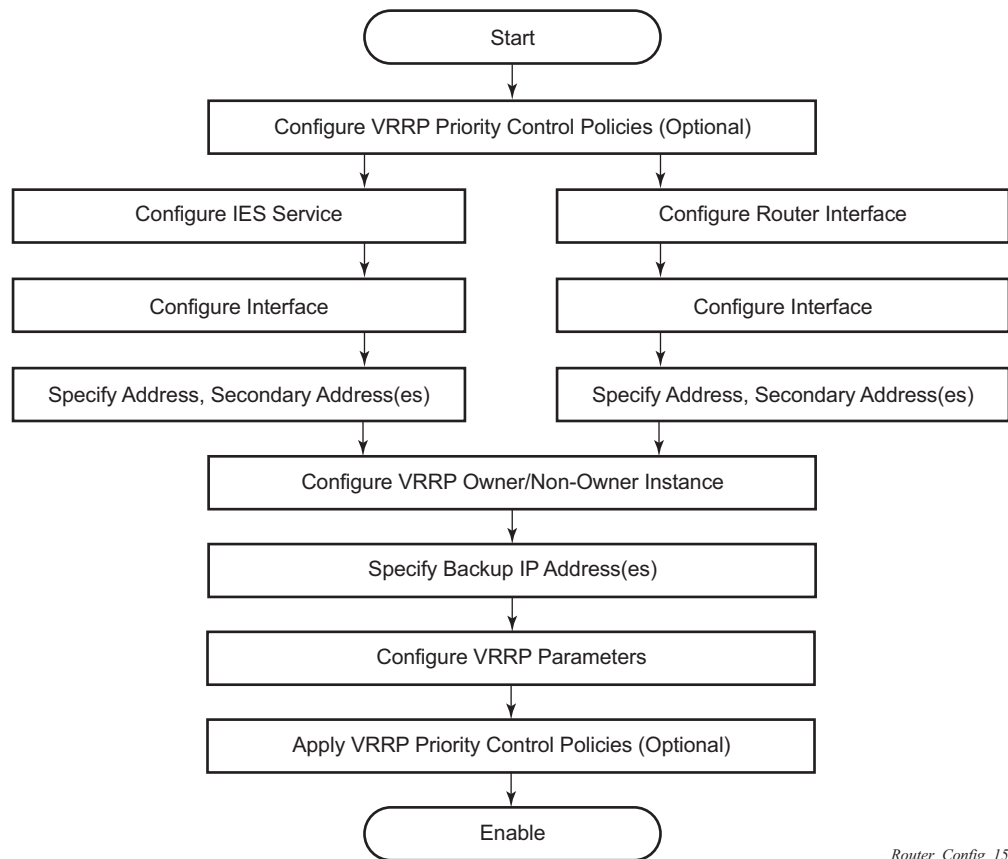
Figure 24 shows part 1 of the process to configure and implement VRRP parameters.

Figure 24 VRRP Configuration and Implementation Flow - Part 1



Router_Config_14

Figure 25 VRRP Configuration and Implementation Flow



Router_Config_15

3.6 Configuration Notes

This section describes VRRP configuration restrictions.

3.6.1 General

- Creating and applying VRRP policies are optional.
- Backup command:
 - The backup IP addresses must be on the same subnet. The backup addresses explicitly define which IP addresses are in the VRRP advertisement message IP address list.
 - In the owner mode, the backup IP address must be identical to one of the interface's IP addresses. The backup address explicitly defines which IP addresses are in the VRRP advertisement message IP address list.
 - For IPv6, one of the backup addresses configured must be the link-local address of the owner VRRP instance.

3.7 Configuring VRRP with CLI

This section provides information to configure VRRP using the command line interface.

3.7.1 VRRP Configuration Overview

Configuring VRRP policies and configuring VRRP instances on interfaces and router interfaces is optional. The basic owner and non-owner VRRP configurations on an IES or router interface must specify the **backup ip-address** parameter.

VRRP helps eliminate the single point of failure in a routed environment by using a virtual router IP address shared between two or more routers connecting the common domain. VRRP provides dynamic failover of the forwarding responsibility if the master becomes unavailable.

The VRRP implementation allows one master per IP subnet. All other VRRP instances in the same domain must be in backup mode.

3.7.1.1 Preconfiguration Requirements

VRRP policies:

- VRRP policies must be configured before they can be applied to an interface or IES/VRPN VRRP instance. VRRP policies are configured in the **config>vrrp** context.

Configuring VRRP on an IES or VRPN service interface:

- The service customer account must be created before configuring an IES or VRPN VRRP instance.
- The interface address must be specified in the both the owner and non-owner IES, VRPN, or router interface instances.

3.7.2 Basic VRRP Configurations

Configure VRRP parameters in the following contexts.

3.7.2.1 VRRP Policy

Configuring and applying VRRP policies are optional. There are no default VRRP policies. Each policy must be explicitly defined. A VRRP configuration must include the following:

- Policy ID
- Define at least one of the following priority events:
 - Port down
 - LAG port down
 - Host unreachable
 - Route unknown

The following example shows a sample configuration of a VRRP policy for the 7450 ESS:

```
A:SR2>config>vrrp>policy# info
-----
      delta-in-use-limit 50
      priority-event
        port-down 4/1/2
          hold-set 43200
          priority 100 delta
        exit
        port-down 4/1/3
          priority 200 explicit
        exit
        lag-port-down 1
          number-down 3
          priority 50 explicit
        exit
        exit
        host-unreachable 10.10.24.4
          drop-count 25
        exit
        route-unknown 10.10.0.0/32
        priority 50 delta
        exit
      exit
-----
```

The following example shows a sample configuration of a VRRP policy for the 7750 SR and 7950 XRS:

```
A:SR2>config>vrrp>policy# info
-----
      delta-in-use-limit 50
      priority-event
        port-down 4/1/2
          hold-set 43200
          priority 100 delta
      exit
-----
```

```

exit
port-down 4/1/3
    priority 200 explicit
exit
lag-port-down 1
    number-down 3
    priority 50 explicit
exit
exit
host-unreachable 10.10.24.4
    drop-count 25
exit
route-unknown 10.10.0.0/32
    priority 50 delta
    protocol bgp
exit
exit
-----

```

3.7.2.2 VRRP IES Service Parameters

VRRP parameters are configured within an IES service with two contexts: owner or non-owner. The status is specified when the VRRP configuration is created. When configured as owner, the virtual router instance owns the backup IP addresses. All other virtual router instances participating in this message domain must have the same VRID configured and cannot be configured as owner.

For IPv4, up to four virtual router IDs can be configured on an IES service interface. Each virtual router instance can manage up to 16 backup IP addresses. For IPv6, only one virtual router instance can be configured on an IES service interface.

VRRP parameters configured within an IES service must include the following:

- VRID
- Backup IP addresses

The following example shows a sample IES service owner and non-owner VRRP configuration:

```

A:SR2>config>service>ies# info
-----
interface "tuesday" create
    address 10.10.36.2/24
    sap 7/1/1.2.2 create
    vrrp 19 owner
        backup 10.10.36.2
        authentication-key "testabc"
    exit
exit
interface "testing" create

```

```

        address 10.10.10.16/24
        sap 1/1/55:0 create
        vrrp 12
            backup 10.10.10.15
            policy 1
            authentication-key "testabc"
        exit
    exit
    no shutdown
-----
A:SR2>config>service>ies#

```

3.7.2.2.1 Configure VRRP for IPv6

The following example shows a VRRP for IPV6 configuration and applies to the 7750 SR and 7950 XRS. The interface must be configured first.

```

*A:nlt7750-3>config>router>router-advert# info
-----
        interface "DSC-101-Application"
            use-virtual-mac
            no shutdown
        exit
...
-----
*A:nlt7750-3>config>router>router-advert#

*A:nlt7750-3>config>service>ies# info
-----
        description "VLAN 921 for DSC-101 Application"
        interface "DSC-101-Application" create
            address 10.152.2.220/28
            vrrp 217
                backup 10.152.2.222
                priority 254
                ping-reply
            exit
        ipv6
            address FD10:D68F:1:221::FFFD/64
            link-local-address FE80::D68F:1:221:FFFD preferred
            vrrp 219
                backup FE80::D68F:1:221:FFFF
                priority 254
                ping-reply
            exit
        exit
        sap ccag-1.a:921 create
            description "cross connect to VPLS 921"
        exit
    exit
    no shutdown
-----
*A:nlt7750-3>config>service>ies#

```

3.7.2.3 VRRP Router Interface Parameters

VRRP parameters are configured on a router interface with two contexts: owner or non-owner. The status is specified when the VRRP configuration is created. When configured as owner, the virtual router instance owns the backed up IP addresses. All other virtual router instances participating in this message domain must have the same VRID configured and cannot be configured as owner.

For IPv4, up to four virtual router IDs can be configured on a router interface. Each virtual router instance can manage up to 16 backup IP addresses. For IPv6, only one virtual router instance can be configured on a router interface.

VRRP parameters configured on a router interface must include the following:

- VRID
- Backup IP addresses

The following example shows a sample router interface owner and non-owner VRRP configuration:

```
A:SR4>config>router# info
#-----
echo "IP Configuration "
#-----
      interface "system"
        address 10.10.0.4/32
      exit
      interface "test1"
        address 10.10.14.1/24
        secondary 10.10.16.1/24
        secondary 10.10.17.1/24
        secondary 10.10.18.1/24
      exit
      interface "test2"
        address 10.10.10.23/24
        vrrp 1 owner
          backup 10.10.10.23
          authentication-key "testabc"
        exit
      exit
#-----
A:SR4>config>router#
```

3.7.3 Common Configuration Tasks

This section provides a brief overview of the tasks that must be performed to configure VRRP and provides the CLI commands.

VRRP parameters are defined under a service interface or a router interface context. An IP address must be assigned to each IP interface. Only one IP address can be associated with an IP interface but several secondary IP addresses also be associated.

Owner and non-owner configurations must include the following parameters:

- All participating routers in a VRRP instance must be configured with the same VRID.
- All participating non-owner routers can specify up to 16 backup IP addresses (IP addresses that the master is representing). The owner configuration must include at least one backup IP address.
- For IPv6, all participating routers must be configured with the same link-local backup address (the one configured for the owner instance).

Other owner and non-owner configurations include the following optional commands:

- authentication-key
- MAC
- message-interval

In addition to the common parameters, the following non-owner commands can be configured:

- master-int-inherit
- priority
- policy
- ping-reply
- preempt
- telnet-reply
- ssh-reply (IPv4 only)
- [no] shutdown

3.7.3.1 Creating Interface Parameters

If you have multiple subnets configured on an Ethernet interface, you can configure VRRP on each subnet.

The following displays an IP interface configuration example:

```
A:SR1>config>router# info
#-----
```

```

echo "IP Configuration "
#-----
    interface "system"
        address 10.10.0.1/32
    exit
    interface "testA"
        address 123.123.123.123/24
    exit
    interface "testB"
        address 10.10.14.1/24
        secondary 10.10.16.1/24
        secondary 10.10.17.1/24
        secondary 10.10.18.1/24
    exit
    router-id 10.10.0.1
#-----
A:SR1>config>router#

```

3.7.4 Configuring VRRP Policy Components

The following displays a VRRP policy configuration example:

```

A:SR1>config>vrrp# info
-----
    policy 1
        delta-in-use-limit 50
        priority-event
            port-down 1/1/2
                hold-set 43200
                priority 100 delta
            exit
        route-unknown 0.0.0.0/0
            protocol isis
        exit
    exit
exit
-----
A:SR1>config>vrrp#

```

3.7.4.1 Configuring Service VRRP Parameters

VRRP parameters can be configured on an interface in a service to provide virtual default router support, which allows traffic to be routed without relying on a single router in case of failure. VRRP can be configured in the following two ways.

3.7.4.1.1 Non-Owner VRRP Example

The following displays a basic non-owner VRRP configuration example:

```
A:SR2>config>service>ies# info
-----
...
        interface "testing" create
            address 10.10.10.16/24
            sap 1/1/55:0 create
            vrrp 12
                backup 10.10.10.15
                policy 1
                authentication-key "testabc"
            exit
        exit
        no shutdown
-----
A:SR2>config>service>ies#
```

3.7.4.1.2 Owner Service VRRP

The following displays an owner service VRRP configuration example:

```
A:SR4>config>router# info
#-----
echo "IP Configuration "
#-----
...
        interface "test2"
            address 10.10.10.23/24
            vrrp 1 owner
                backup 10.10.10.23
                authentication-key "testabc"
            exit
        exit
#-----
A:SR4>config>router#
```

3.7.4.2 Configuring Router Interface VRRP Parameters

VRRP parameters can be configured on an interface in an interface to provide virtual default router support, which allows traffic to be routed without relying on a single router in case of failure. VRRP can be configured in following two ways.

3.7.4.2.1 Router Interface VRRP Non-Owner

The following displays a router interface non-owner VRRP configuration example:

```
A:SR2>config># info
#-----
    interface "if-test"
        address 10.20.30.40/24
        secondary 10.10.50.1/24
        secondary 10.10.60.1/24
        secondary 10.10.70.1/24
        vrrp 1
            backup 10.10.50.2
            backup 10.10.60.2
            backup 10.10.70.2
            backup 10.20.30.41
            ping-reply
            telnet-reply
            authentication-key "testabc"
        exit
    exit
#-----
A:SR2>config>#
```

3.7.4.2.2 Router Interface VRRP Owner

The following displays a router interface owner VRRP configuration example:

```
A:SR2>config>router# info
#-----
    interface "vrrpowner"
        address 10.10.10.23/24
        vrrp 1 owner
            backup 10.10.10.23
            authentication-key "testabc"
        exit
    exit
#-----
A:SR2>config>router#
```

3.8 VRRP Configuration Management Tasks

This section describes VRRP configuration management tasks:

3.8.1 Modifying a VRRP Policy

To access a specific VRRP policy, you must specify the policy ID. To display a list of VRRP policies, use the `show vrrp policy` command.

The following example shows the modified VRRP policy configuration:

```
A:SR2>config>vrrp>policy# info
-----
      delta-in-use-limit 50
      priority-event
        port-down 1/1/2
          hold-set 43200
          priority 100 delta
        exit
      port-down 1/1/3
        priority 200 explicit
      exit
      host-unreachable 10.10.24.4
        drop-count 25
      exit
    exit
-----
A:SR2>config>vrrp>policy#
```

3.8.1.1 Deleting a VRRP Policy

Policies are only applied to non-owner VRRP instances. A VRRP policy cannot be deleted if it is applied to an interface or to an IES service. Each instance in which the policy is applied must be deleted.

The Applied column in the following example shows whether the VRRP policies are applied to an entity.

```
A:SR2#
=====
VRRP Policies
=====
```

| Policy Id | Current Priority & Effect | Current Explicit | Current Delta Sum | Delta Limit | Applied |
|-----------|---------------------------|------------------|-------------------|-------------|---------|
| 1 | 200 Explicit | 200 | 100 | 50 | Yes |
| 15 | 254 | None | None | 1 | No |

```

32          100          None          None          1          No
=====
A:SR2#

```

3.8.2 Modifying Service and Interface VRRP Parameters

3.8.2.1 Modifying Non-Owner Parameters

After a VRRP instance is created as non-owner, it cannot be modified to the owner state. The VRID must be deleted, then recreated with the owner keyword, to invoke IP address ownership.

3.8.2.2 Modifying Owner Parameters

After a VRRP instance is created as owner, it cannot be modified to the non-owner state. The VRID must be deleted, then recreated without the owner keyword, to remove IP address ownership.

Entering the owner keyword is optional when entering the VRID for modification purposes.

3.8.2.3 Deleting VRRP from an Interface or Service

The VRID does not need to be shutdown to remove the virtual router instance from an interface or service.

Example:

```

config>router#interface
config>router# interface if-test
config>router>if# shutdown
config>router>if# exit
config>router# no interface if-test
config>router#

```

The following example shows the command usage to delete a VRRP instance from an interface or IES service:

Example:

```

config>service#ies 10
config>service>ies# interface "test"
config>service>ies>if# vrrp 1

```

```
config>service>ies>if>vrrp# shutdown
config>service>ies>if>vrrp# exit
config>service>ies>if# no vrrp 1
config>service>ies>if# exit all
```

3.9 VRRP Configuration Command Reference

- [Command Hierarchies](#)
- [Command Descriptions](#)

3.9.1 Command Hierarchies

- [IPv4 Interface VRRP Commands](#)
- [Router Interface Commands](#)
- [IPv6 Interface VRRP Commands](#)
- [Priority Control Event Policy Commands](#)

3.9.1.1 IPv4 Interface VRRP Commands

```

config
  — router
    — [no] interface interface-name
      — vrrp virtual-router-id [owner] [passive]
      — no vrrp virtual-router-id
        — authentication-key [authentication-key | hash-key] [hash | hash2]
        — no authentication-key
        — [no] backup ip-address
        — [no] bfd-enable service-id interface interface-name dst-ip ip-address
        — [no] bfd-enable interface interface-name dst-ip ip-address
        — init-delay seconds
        — no init-delay
        — mac mac-address
        — no mac
        — [no] master-int-inherit
        — message-interval {[seconds] [milliseconds milliseconds]}
        — no message-interval
        — [no] ping-reply
        — oper-group group-name
        — no oper-group
        — policy policy-id
        — no policy
        — [no] preempt
        — priority priority
        — no priority
        — [no] ssh-reply
        — [no] standby-forwarding
        — [no] telnet-reply
        — [no] shutdown

```

- [no] **traceroute-reply**

VRRP commands are applicable to router interfaces, IES interfaces and VPRN. The **authentication-key**, **bfd-enable**, and **ssh-reply** commands are applicable only to IPv4 contexts, not IPv6.

3.9.1.2 Router Interface Commands

```

config
— router [router-name]
  — [no] interface ip-int-name
    — [no] ipv6
      — address ipv6-address/prefix-length [eui-64]
      — no address ipv6-address/prefix-length
      — icmp6
        — packet-too-big [number seconds]
        — no packet-too-big
        — param-problem [number seconds]
        — no param-problem
        — redirects [number seconds]
        — no redirects
        — time-exceeded [number seconds]
        — no time-exceeded
        — unreachables [number seconds]
        — no unreachables
      — link-local-address ipv6-address [preferred]
      — no link-local-address
      — [no] local-proxy-nd
      — neighbor ipv6-address [mac-address]
      — no neighbor ipv6-address
      — proxy-nd-policy policy-name [policy-name...(up to 5 max)]
      — no proxy-nd-policy

```

3.9.1.3 IPv6 Interface VRRP Commands

The IPv6 interface commands only apply to the 7750 SR and 7950 XRS.

```

config
— router [router-name]
  — [no] interface ip-int-name
    — [no] ipv6
      — vrrp virtual-router-id [owner] [passive]
      — no vrrp virtual-router-id
        — [no] backup ipv6-address
        — [no] bfd-enable service-id interface interface-name dst-ip ip-address
        — [no] bfd-enable interface interface-name dst-ip ip-address

```

- **init-delay** *seconds*
- **no init-delay**
- **mac** *mac-address*
- **no mac**
- **[no] master-int-inherit**
- **message-interval** *{[seconds] [milliseconds milliseconds]}*
- **no message-interval**
- **[no] ping-reply**
- **policy** *vrrp-policy-id*
- **no policy**
- **[no] preempt**
- **priority** *priority*
- **no priority**
- **[no] shutdown**
- **[no] standby-forwarding**
- **[no] telnet-reply**
- **[no] traceroute-reply**

3.9.1.4 Priority Control Event Policy Commands

- ```

config
 — vrrp
 — [no] policy policy-id [context service-id
 — delta-in-use-limit limit
 — no delta-in-use-limit
 — description description string
 — no description
 — [no] priority-event
 — [no] host-unreachable ip-address
 — [no] host-unreachable ipv6-address
 — drop-count consecutive-failures
 — no drop-count
 — hold-clear seconds
 — no hold-clear
 — hold-set seconds
 — no hold-set
 — interval seconds
 — no interval
 — padding-size size
 — no padding-size
 — priority priority-level [{delta | explicit}]
 — no priority
 — timeout seconds
 — no timeout
 — [no] lag-port-down lag-id
 — hold-clear seconds
 — no hold-clear
 — hold-set seconds
 — no hold-set
 — [no] number-down number-of-lag-ports-down
 — priority priority-level [{delta | explicit}]

```

---

```

 — no priority
 — weight-down lag-ports-down-weight
 — no weight-down
— mc-ipsec-non-forwarding tunnel-grp-id
 — hold-clear seconds
 — no hold-clear
 — hold-set seconds
 — no hold-set
 — priority priority-level [{delta | explicit}]
 — no priority
— [no] port-down port-id
 — hold-clear seconds
 — no hold-clear
 — hold-set seconds
 — no hold-set
 — priority priority-level [{delta | explicit}]
 — no priority
— [no] route-unknown ip-prefix/mask
 — hold-clear seconds
 — no hold-clear
 — hold-set seconds
 — no hold-set
 — less-specific [allow-default]
 — no less-specific
 — [no] next-hop ip-address
 — priority priority-level [delta | explicit]
 — no priority
 — protocol protocol
 — no protocol [protocol]
 — [no] protocol {bgp | bgp -vpn | ospf | isis | rip | static}
— [no] shutdown

```

### 3.9.2 Command Descriptions

- [Interface Configuration Commands](#)
- [Priority Policy Commands](#)
- [Priority Policy Event Commands](#)
- [Priority Policy Port Down Event Commands](#)
- [Priority Policy LAG Events Commands](#)
- [Priority Policy Host Unreachable Event Commands](#)
- [Priority Policy Route Unknown Event Commands](#)



### 3.9.2.1 Interface Configuration Commands

#### authentication-key

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>authentication-key</b> [ <i>authentication-key</i>   <i>hash-key</i> ] [ <b>hash</b>   <b>hash2</b> ]<br><b>no authentication-key</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>router>if>vrrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command sets the simple text authentication key used to generate master VRRP advertisement messages and validates VRRP advertisements.</p> <p>If simple text password authentication is not required, the <b>authentication-key</b> command is not required.</p> <p>The command is configurable in both non-owner and owner <b>vrrp</b> nodal contexts.</p> <p>The <i>key</i> parameter identifies the simple text password to be used when VRRP Authentication Type 1 is enabled on the virtual router instance. Type 1 uses an eight octet long string that is inserted into all transmitted VRRP advertisement messages and is compared against all received VRRP advertisement messages. The authentication data fields are used to transmit the <i>key</i>.</p> <p>The <i>key</i> string is case sensitive and is left justified in the VRRP advertisement message authentication data fields. The first field contains the first four characters with the first octet (starting with IETF RFC bit position 0) containing the first character. The second field similarly holds the fifth through eighth characters. Any unspecified portion of the authentication data field is padded with a 0 value in the corresponding octet.</p> <p>If the command is re-executed with a different password key defined, the new key is used immediately.</p> <p>The <b>authentication-key</b> command can be executed at anytime.</p> <p>To change the current in-use password key on multiple virtual router instances:</p> <p>Identify the current master.</p> <ol style="list-style-type: none"> <li>1. Shutdown the virtual router instance on all backups.</li> <li>2. Execute the <b>authentication-key</b> command on the master to change the password key.</li> <li>3. Execute the <b>authentication-key</b> command and <b>no shutdown</b> command on each backup.</li> </ol> <p>The <b>no</b> form of the command reverts to the default value.</p> |
| <b>Default</b>     | no authentication-key — The authentication key value is the null string.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

- Parameters**
- authentication-key* — The authentication key. Allowed values are any string up to 8 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.
  - hash-key* — The hash key. The key can be any combination of ASCII characters up to 22 (hash-key1) or 121 (hash-key2) characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (" ").  
This is useful when a user must configure the parameter, but for security purposes, the actual unencrypted key value is not provided.
  - hash** — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified
  - hash2** — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

## backup

- Syntax** [no] **backup** *ip-address*
- Context** config>router>if>vrrp
- Description** This command associates router IP addresses with the parental IP interface IP addresses.
- The **backup** command has two distinct functions when used in an **owner** or a **non-owner** context of the virtual router instance.
- Non-owner virtual router instances actually create a routable IP interface address that is operationally dependent on the virtual router instance mode (master or backup). The **backup** command in **owner** virtual router instances does not create a routable IP interface address; it simply defines the existing parental IP interface IP addresses that are advertised by the virtual router instance.
- For **owner** virtual router instances, the **backup** command defines the IP addresses that are advertised within VRRP advertisement messages. This communicates the IP addresses that the master is representing to backup virtual routers receiving the messages. Advertising a correct list is important. The specified *ip-addr* must be equal to one of the existing parental IP interface IP addresses (primary or secondary) or the **backup** command will fail.

For non-owner virtual router instances, the **backup** command actually creates an IP interface IP address used for routing IP packets and communicating with the system when the access commands are defined (**ping-reply**, **telnet-reply**, and **ssh-reply**). The specified *ip-addr* must be an IP address that is within one of the parental IP interface local subnets created with the **address** or **secondary** commands. If a local subnet does not exist that includes the specified *ip-addr* or if *ip-addr* is the same IP address as the parental IP interface IP address, the **backup** command will fail.

The new interface IP address created with the **backup** command assumes the mask and parameters of the corresponding parent IP interface IP address. The *ip-addr* is only active when the virtual router instance is operating in the master state. When not operating as master, the virtual router instance acts as if it is operationally down. It will not respond to ARP requests to *ip-addr*, nor will it route packets received with its *vrid* derived source MAC address. A non-master virtual router instance always silently discards packets destined to *ip-addr*. A single virtual router instance may only have a single virtual router IP address from a given parental local subnet. Multiple virtual router instances can define a virtual router IP address from the same local subnet as long as each is a different IP address.

In IPv4, up to sixteen **backup ip-addr** commands can be executed within the same virtual router instance. Executing **backup** multiple times with the same *ip-addr* results in no operation performed and no error generated. At least one successful **backup ip-addr** command must be executed before the virtual router instance can enter the operational state.

When operating as (non-owner) master, the default functionality associated with *ip-addr* is ARP response to ARP requests to *ip-addr*, routing of packets destined to the virtual router instance source MAC address and silently discarding packets destined to *ip-addr*. Enabling the non-owner-access parameters selectively allows ping, Telnet and SSH connectivity to *ip-addr* when the virtual router instance is operating as master.

The **no** form of the command removes the specified virtual router IP address from the virtual router instance. For non-owner virtual router instances, this causes all routing and local access associated with the *ip-addr* to cease. For **owner** virtual router instances, the **no backup** command only removes *ip-addr* from the list of advertised IP addresses. If the last *ip-addr* is removed from the virtual router instance, the virtual router instance will enter the operationally down state

|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>       | no backup — No virtual router IP address is assigned.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Special Cases</b> | <p><b>Assigning the Virtual Router ID IP Address</b> — Once the <i>vrid</i> is created on the parent IP interface, IP addresses need to be assigned to the virtual router instance. If the <i>vrid</i> was created with the keyword <b>owner</b>, the virtual router instance IP addresses must have one or more of the parent IP interface defined IP addresses (primary and secondary). For non-owner virtual router instances, the virtual router IP addresses each must be within one of the parental IP interface IP address defined local subnets. For both <b>owner</b> and non-owner virtual router instances, the virtual router IP addresses must be explicitly defined using the <b>backup ip-addr</b> command.</p> |

**Virtual Router Instance IP Address Assignment Conditions** — The RFC does not specify that the assigned IP addresses to the virtual router instance must be in the same subnet as the parent IP interface primary IP address or secondary IP addresses. The only requirement is that all virtual routers participating in the same virtual router instance have the same virtual router IP addresses assigned. To avoid confusion, the assigned virtual router IP addresses must be in a local subnet of one of the parent IP interfaces IP addresses. For **owner** virtual router instances the assigned virtual router IP address must be the same as one of the parental IP interface primary or secondary IP addresses.

The following rules apply when adding, changing, or removing parental and virtual router IP addresses:

**Owner Virtual Router IP Address Parental Association** — When an IP address is assigned to an **owner** virtual router instance, it must be associated with one of the parental IP interface-assigned IP addresses. The virtual router IP address must be equal to the primary or one of the secondary IP addresses within the parental IP interface.

**Table 36 Example - Owner Virtual Router Instance**

|                              |                                  |                                                    |
|------------------------------|----------------------------------|----------------------------------------------------|
| Parent IP addresses:         | 10.10.10.10/24<br>11.11.11.11/24 |                                                    |
| Virtual router IP addresses: | 10.10.10.11                      | Invalid (not equal to parent IP address)           |
|                              | 10.10.10.10                      | Associated (same as parent IP address 10.10.10.10) |
|                              | 10.10.11.11                      | Invalid (not equal to parent IP address)           |
|                              | 11.11.11.254                     | Invalid (not equal to parent IP address)           |
|                              | 11.11.11.255                     | Invalid (not equal to parent IP address)           |

**Non-Owner Virtual Router IP Address Parental Association** — When an IP address is assigned to a non-owner virtual router instance, it must be associated with one of the parental IP interface assigned IP addresses. The virtual router IP address must be a valid IP address within one of the parental IP interfaces local subnet. Local subnets are created by the primary or secondary IP addresses in conjunction with the IP addresses mask. If the defined virtual router IP address is equal to the associated subnet's broadcast address, it is invalid. Virtual router IP addresses for non-owner virtual router instances that are equal to a parental IP interface IP address are also invalid.

The same virtual router IP address may not be assigned to two separate virtual router instances. If the virtual router IP address already exists on another virtual router instance, the virtual router IP address assignment will fail.

**Table 37 Example - Non-Owner Virtual Router Instance**

|                              |                                  |                                               |
|------------------------------|----------------------------------|-----------------------------------------------|
| Parent IP addresses:         | 10.10.10.10/24<br>11.11.11.11/24 |                                               |
| Virtual router IP addresses: | 10.10.10.11                      | Associated with 10.10.10.10 (in subnet)       |
|                              | 10.10.10.10                      | Invalid (same as parent IP address)           |
|                              | 10.10.11.11                      | Invalid (outside of all Parent IP subnets)    |
|                              | 11.11.11.254                     | Associated with 11.11.11.11 (in subnet)       |
|                              | 11.11.11.255                     | Invalid (broadcast address of 11.11.11.11/24) |

**Virtual Router IP Address Assignment without Parent IP Address** — When assigning an IP address to a virtual router instance, an associated IP address (see [backup](#) and [backup](#)) on the parental IP interface must already exist. If an associated IP address on the parental IP interface is not configured, the virtual router IP address assignment fails.

**Parent Primary IP Address Changed** — When a virtual router IP address is set and the associated parent IP interface IP address is changed, the new parent IP interface IP address is evaluated to ensure it meets the association rules defined in [backup](#) or [backup](#). If the association check fails, the parental IP address change is not allowed. If the parental IP address change fails, the previously configured IP address definition remains in effect.

Only the primary parent IP address can be changed. Secondary addresses must be removed before the new IP address can be added. [backup](#) explains IP address removal conditions.

**Parent Primary or Secondary IP Address Removal** — When a virtual router IP address is successfully set, but removing the associated parent IP interface IP address is attempted and fails. All virtual router IP addresses associated with the parental IP interface IP address must be deleted prior to removing the parental IP address. This includes virtual router IP address associations from multiple virtual router instances on the IP interface.

**Parameters** *ip-address* — The virtual router IP address expressed in dotted decimal notation. The IP virtual router IP address must be in the same subnet of the parental IP interface IP address or equal to one of the primary or secondary IP addresses for **owner** virtual router instances.

**Values** 1.0.0.1 - 223.255.255.254

---

## backup

|                    |                                                                                            |
|--------------------|--------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>backup</b> <i>ipv6-address</i><br><b>no backup</b>                                      |
| <b>Context</b>     | config>router>if>ipv6>vrrp                                                                 |
| <b>Description</b> | This command associates router IPv6 addresses with the parental IP interface IP addresses. |

The **backup** command has two distinct functions when used in an **owner** or a **non-owner** context of the virtual router instance.

Non-owner virtual router instances actually create a routable IP interface address that is operationally dependent on the virtual router instance mode (master or backup). The **backup** command in **owner** virtual router instances does not create a routable IP interface address; it simply defines the existing parental IP interface IP addresses that are advertised by the virtual router instance.

For **owner** virtual router instances, the **backup** command defines the IP addresses that are advertised within VRRP advertisement messages. This communicates the IP addresses that the master is representing to backup virtual routers receiving the messages. Advertising a correct list is important. The specified *ipv6-addr* must be equal to one of the existing parental IP interface IP addresses (link-local or global) or the **backup** command will fail.

For non-owner virtual router instances, the **backup** command actually creates an IP interface IP address used for routing IP packets and communicating with the system when the access commands are defined (**ping-reply**, **telnet-reply**, and **ssh-reply**). The specified *ipv6-addr* must be an IP address that is within one of the parental IP interface local subnets created with the **link-local-address** or **address** commands. If a local subnet does not exist that includes the specified *ipv6-addr* or if *ipv6-addr* is the same IP address as the parental IP interface IP address, the **backup** command will fail.

The new interface IP address created with the **backup** command assumes the mask and parameters of the corresponding parent IP interface IP address. The *ipv6-addr* is only active when the virtual router instance is operating in the master state. For IPv6 VRRP, the parental interface's IP address that is in the same subnet as the backup address must be manually-configured, non EUI-64 and configured to be in the preferred state.

When not operating as master, the virtual router instance acts as if it is operationally down. It will not respond to ARP requests to *ipv6-addr*, nor will it route packets received with its *vrid* derived source MAC address. A non-master virtual router instance always silently discards packets destined to *ipv6-addr*. A single virtual router instance may only have a single virtual router IP address from a specific parental local subnet. Multiple virtual router instances can define a virtual router IP address from the same local subnet as long as each is a different IP address.

Executing **backup** multiple times with the same *ipv6-addr* results in no operation performed and no error generated. At least one successful **backup** *ipv6-addr* command must be executed before the virtual router instance can enter the operational state.

When operating as (non-owner) master, the default functionality associated with *ipv6-addr* is ARP response to ARP requests to *ip-addr*, routing of packets destined to the virtual router instance source MAC address and silently discarding packets destined to *ipv6-addr*. An IPv6 virtual router instance can enter the operational state only if one of the configured backup address is a link-local address and the router advertisement of the interface is configured to use the virtual MAC address. Enabling the non-owner-access parameters selectively allows ping, Telnet and traceroute connectivity to *ipv6-addr* when the virtual router instance is operating as master.

The **no** form of the command removes the specified virtual router IP address from the virtual router instance. For non-owner virtual router instances, this causes all routing and local access associated with the *ipv6-addr* to cease. For **owner** virtual router instances, the **no backup** command only removes *ipv6-addr* from the list of advertised IP addresses. If the last *ipv6-addr* or the *link-local* address is removed from the virtual router instance, the virtual router instance will enter the operationally down state

**Default** no backup — No virtual router IP address is assigned.

**Special Cases** **Assigning the Virtual Router ID Address** — Once the *vrid* is created on the parent IP interface, IP addresses need to be assigned to the virtual router instance. If the *vrid* was created with the keyword **owner**, the virtual router instance IP addresses must have one or more of the parent IP interface defined IP addresses. For non-owner virtual router instances, the virtual router IP addresses each must be within one of the parental IP interface IP address defined local subnets. For both **owner** and non-owner virtual router instances, the virtual router IP addresses must be explicitly defined using the **backup** *ipv6-addr* command.

The following rules apply when adding, changing, or removing parental and virtual router IP addresses:

**Owner Virtual Router IP Address Parental Association** — When an IP address is assigned to an **owner** virtual router instance, it must be associated with one of the parental IP interface-assigned IP addresses.

**Table 38 Example - Owner Virtual Router Instance**

|                              |                                  |                                                    |
|------------------------------|----------------------------------|----------------------------------------------------|
| Parent IP addresses:         | 10.10.10.10/24<br>11.11.11.11/24 |                                                    |
| Virtual router IP addresses: | 10.10.10.11                      | Invalid (not equal to parent IP address)           |
|                              | 10.10.10.10                      | Associated (same as parent IP address 10.10.10.10) |
|                              | 10.10.11.11                      | Invalid (not equal to parent IP address)           |
|                              | 11.11.11.254                     | Invalid (not equal to parent IP address)           |
|                              | 11.11.11.255                     | Invalid (not equal to parent IP address)           |

**Non-Owner Virtual Router IP Address Parental Association** — When an IP address is assigned to a non-owner virtual router instance, it must be associated with one of the parental IP interface assigned IP addresses. The virtual router IP address must be a valid IP address within one of the parental IP interfaces local subnet. Local subnets are created by the link-local or global IP addresses in conjunction with the IP addresses mask. If the defined virtual router IP address is equal to the associated subnet's broadcast address, it is invalid. Virtual router IP addresses for non-owner virtual router instances that are equal to a parental IP interface IP address are also invalid.

The same virtual router IP address may not be assigned to two separate virtual router instances. If the virtual router IP address already exists on another virtual router instance, the virtual router IP address assignment will fail.

One exception to this rule is for the IPv6 link-local address that is configured as a backup address. The same link-local address can be configured in all virtual routers that use the same vrid.

**Table 39      Example - Non-Owner Virtual Router Instance**

|                                |                                  |                                               |
|--------------------------------|----------------------------------|-----------------------------------------------|
| Parent IP addresses:           | 10.10.10.10/24<br>11.11.11.11/24 |                                               |
| Virtual router IPv6 addresses: | 10.10.10.11                      | Associated with 10.10.10.10 (in subnet)       |
|                                | 10.10.10.10                      | Invalid (same as parent IP address)           |
|                                | 10.10.11.11                      | Invalid (outside of all Parent IP subnets)    |
|                                | 11.11.11.254                     | Associated with 11.11.11.11 (in subnet)       |
|                                | 11.11.11.255                     | Invalid (broadcast address of 11.11.11.11/24) |

**Virtual Router IP Address Assignment without Parent IP Address** — When assigning an IP address to a virtual router instance, an associated IP address (see [backup](#) and [backup](#)) on the parental IP interface must already exist. If an associated IP address on the parental IP interface is not configured, the virtual router IP address assignment fails.

**Virtual Router IPv6 Address Assignment** — An IPv6 backup address requires that the parental IP address that is in the same subnet as the backup address must be manually configured, non-EUI-64 and configured to be in the preferred state.

**Parameters**    *ipv6-address* — The virtual router IP address expressed in dotted decimal notation. The IP virtual router IP address must be in the same subnet of the parental IP interface IP address or equal to one of the parent interface addresses for **owner** virtual router instances.

**Values**



|                  |                                                                                                      |
|------------------|------------------------------------------------------------------------------------------------------|
| ipv6-<br>address | x::x::x::x::x::x (eight 16-bit<br>pieces)<br>x::x::x::x::x::d.d.d.d<br>x: [0..FFFF]H<br>d: [0..255]D |
|------------------|------------------------------------------------------------------------------------------------------|

## bfd-enable

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] bfd-enable</b> [ <i>service-id</i> ] <b>interface</b> <i>interface-name</i> <b>dst-ip</b> <i>ip-address</i><br><b>[no] bfd-enable interface</b> <i>interface-name</i> <b>dst-ip</b> <i>ip-address</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | config>router>if>vrrp<br>config>router>if>ipv6>vrrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This commands assigns a bidirectional forwarding detect (BFD) session to a specific VRRP/SRRP instance. This BFD sessions provided a heartbeat mechanism that can be used to speed up the transition of the standby VRRP router to an active state. If the associated BFD session fails, the VRRP routers will immediately send a VRRP Advertisement message. In addition, the standby VRRP router(s) will transition to a Master state to speed convergence. The normal VRRP election process will then take place based on the Advertisement messages sent by all VRRP routers.</p> <p>There can be only one BFD session assigned to any given VRRP/SRRP instance, but there can be multiple SRRP/VRRP sessions using the same BFD session.</p> <p>The parameters used for the BFD sessions are set by the BFD command under the IP interface.</p> <p>The <b>no</b> form of this command removes BFD from the configuration.</p> |
| <b>Parameters</b>  | <p><i>service-id</i> — Specifies the service ID of the interface running BFD.</p> <p><b>Values</b>     <i>service-id</i>: 1 to 2147483647<br/>                  <i>svc-name</i>: 64 characters maximum</p> <p><b>interface</b> <i>interface-name</i> — Specifies the name of the interface running BFD. The specified interface may not yet be configured with BFD. However, when it is, this virtual router will then initiate the BFD session.</p> <p><b>dst-ip</b> <i>ip-address</i> — Specifies the destination address to be used for the BFD session.</p>                                                                                                                                                                                                                                                                                                                                                                       |

## init-delay

|                |                                                          |
|----------------|----------------------------------------------------------|
| <b>Syntax</b>  | <b>init-delay</b> <i>seconds</i><br><b>no init-delay</b> |
| <b>Context</b> | config>router>if>vrrp                                    |

---

```
config>router>if>ipv6>vrrp
```

|                    |                                                                                 |
|--------------------|---------------------------------------------------------------------------------|
| <b>Description</b> | This command configures a VRRP initialization delay timer.                      |
| <b>Default</b>     | no init-delay                                                                   |
| <b>Parameters</b>  | <i>seconds</i> — Specifies the initialization delay timer for VRRP, in seconds. |
| <b>Values</b>      | 1 to 65535                                                                      |

## mac

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>mac</b> <i>mac-address</i><br><b>no mac</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>router>if>vrrp<br>config>router>if>ipv6>vrrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command sets an explicit MAC address used by the virtual router instance overriding the VRRP default derived from the VRID.</p> <p>Changing the default MAC address is useful when an existing HSRP or other non-VRRP default MAC is in use by the IP hosts using the virtual router IP address. Many hosts do not monitor unessential ARPs and continue to use the cached non-VRRP MAC address after the virtual router becomes master of the host's gateway address.</p> <p>The <b>mac</b> command sets the MAC address used in ARP responses when the virtual router instance is master. Routing of IP packets with <i>mac-address</i> as the destination MAC is also enabled. The <b>mac</b> setting must be the same for all virtual routers participating as a virtual router or indeterminate connectivity by the attached IP hosts will result. All VRRP advertisement messages are transmitted with <i>mac-address</i> as the source MAC.</p> <p>The command can be configured in both non-owner and owner <b>vrrp</b> nodal contexts.</p> <p>The <b>mac</b> command can be executed at any time and takes effect immediately. When the virtual router MAC on a master virtual router instance changes, a gratuitous ARP is immediately sent with a VRRP advertisement message. If the virtual router instance is disabled or operating as backup, the gratuitous ARP and VRRP advertisement message is not sent.</p> <p>The <b>no</b> form of the command restores the default VRRP MAC address to the virtual router instance.</p> |
| <b>Default</b>     | no mac — The virtual router instance uses the default VRRP MAC address derived from the VRID.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <i>mac-address</i> — The 48-bit MAC address for the virtual router instance in the form <i>aa:bb:cc:dd:ee:ff</i> or <i>aa-bb-cc-dd-ee-ff</i> where <i>aa</i> , <i>bb</i> , <i>cc</i> , <i>dd</i> , <i>ee</i> and <i>ff</i> are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC, and non-IEEE reserved MAC addresses.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

## master-int-inherit

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] master-int-inherit</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>router>if>vrrp<br>config>router>if>ipv6>vrrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>This command enables the virtual router instance to inherit the master VRRP router's advertisement interval timer which is used by backup routers to calculate the master down timer.</p> <p>The <b>master-int-inherit</b> command is only available in the non-owner nodal context and is used to allow the current virtual router instance master to dictate the master down timer for all backup virtual routers. The <b>master-int-inherit</b> command has no effect when the virtual router instance is operating as master.</p> <p>If <b>master-int-inherit</b> is not enabled, the locally configured <b>message-interval</b> must match the master's VRRP advertisement message advertisement interval field value or the message is discarded.</p> <p>The <b>no</b> form of the command restores the default operating condition which requires the locally configured <b>message-interval</b> to match the received VRRP advertisement message advertisement interval field value.</p> |
| <b>Default</b>     | no master-int-inherit — The virtual router instance does not inherit the master VRRP router's advertisement interval timer and uses the locally configured message interval.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## message-interval

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>message-interval</b> {[seconds] [milliseconds milliseconds]}<br><b>no message-interval</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>router>if>vrrp<br>config>router>if>ipv6>vrrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | <p>This command configures the administrative advertisement message timer used by the master virtual router instance to send VRRP advertisement messages and to derive the master down timer as backup.</p> <p>For an owner virtual router instance, the administrative advertisement timer directly sets the operational advertisement timer and indirectly sets the master down timer for the virtual router instance.</p> <p>Non-owner virtual router instances usage of the <b>message-interval</b> setting is dependent on the state of the virtual router (master or backup) and the state of the <b>master-int-inherit</b> parameter.</p> |

- When a non-owner is operating as master for the virtual router, the configured **message-interval** is used as the operational advertisement timer similar to an owner virtual router instance. The **master-int-inherit** command has no effect when operating as master.
- When a non-owner is in the backup state with **master-int-inherit** disabled, the configured **message-interval** value is used to match the incoming VRRP advertisement message advertisement interval field. If the locally configured message interval does not match the advertisement interval field, the VRRP advertisement is discarded.
- When a non-owner is in the backup state with **master-int-inherit** enabled, the configured **message-interval** is ignored. The master down timer is indirectly derived from the incoming VRRP advertisement message advertisement interval field value.

VRRP advertisements messages that are fragmented, or contain IP options (IPv4), or contain extension headers (IPv6) require a longer message interval to be configured.

The in-use value of the message interval is used to derive the master down timer to be used when the virtual router is operating in backup mode based on the following formula:

$$(3 \times (\text{in-use message interval}) + \text{skew time})$$

The skew time portion is used to slow down virtual routers with relatively low priority values when competing in the master election process.

The command is available in both non-owner and owner **vrrp** nodal contexts.

By default, a **message-interval** of 1 second is used.

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

|            |                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Default    | message-interval 1 — Advertisement timer set to 1 second.                                                                                                                                                                                                                                                                                                                                                                                            |
| Parameters | <p><i>seconds</i> — The number of seconds that will transpire before the advertisement timer expires expressed as a decimal integer.</p> <p><b>Values</b> IPv4: 1 to 255<br/>IPv6: 1 to 40</p> <p><b>milliseconds</b> <i>milliseconds</i> — Specifies the time interval, in milliseconds, between sending advertisement messages. This parameter is not supported on the 7450 ESS-1 chassis.</p> <p><b>Values</b> 100 to 900<br/>IPv6: 10 to 990</p> |

oper-group

|         |                                                             |
|---------|-------------------------------------------------------------|
| Syntax  | <b>oper-group</b> <i>group-name</i><br><b>no oper-group</b> |
| Context | config>router>if>vrrp                                       |

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command configures VRRP to associate with an operational group. When associated, VRRP notifies the operational group of its state changes so that other protocols can monitor it to provide a redundancy mechanism. When VRRP is the master router (MR), the operational group is up; the operational group is down for all other VRRP states.</p> <p>The <b>no</b> form of the command removes the association.</p> |
| <b>Default</b>     | no oper-group — No operational group is configured.                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <i>group-name</i> — Specifies the operational group identifier up to 32 characters in length.                                                                                                                                                                                                                                                                                                                                |

## policy

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>policy</b> <i>policy-id</i><br><b>no policy</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>router>if>vrrp<br>config>router>if>ipv6>vrrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>This command adds a VRRP priority control policy association with the virtual router instance.</p> <p>To further augment the virtual router instance base priority, VRRP priority control policies can be used to override or adjust the base priority value depending on events or conditions within the chassis.</p> <p>The policy can be associated with more than one virtual router instance. The priority events within the policy either override or diminish the base priority set with the <b>priority</b> command dynamically affecting the in-use priority. As priority events clear in the policy, the in-use priority can eventually be restored to the base <b>priority</b> value.</p> <p>The <b>policy</b> command is only available in the non-owner <b>vrrp</b> nodal context. The priority of <b>owner</b> virtual router instances is permanently set to 255 and cannot be changed by VRRP priority control policies. For non-owner virtual router instances, if the <b>policy</b> command is not executed, the base <b>priority</b> is used as the in-use priority.</p> <p>The <b>no</b> form of the command removes existing VRRP priority control policy associations from the virtual router instance. All associations must be removed prior to deleting the policy from the system.</p> |
| <b>Default</b>     | no policy — No VRRP priority control policy is associated with the virtual router instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <i>policy-id</i> — The policy ID of the VRRP priority control expressed as a decimal integer.<br>The <i>vrrp-policy-id</i> must already exist for the command to function.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Values</b>      | 1 to 9999                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

---

## preempt

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] preempt</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>router>if>vrrp<br>config>router>if>ipv6>vrrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>The preempt mode value controls whether a specific backup virtual router preempts a lower priority master.</p> <p>When preempt is enabled, the virtual router instance overrides any non-owner master with an "in use" message priority value less than the virtual router instance in-use priority value. If preempt is disabled, the virtual router only becomes master if the master down timer expires before a VRRP advertisement message is received from another virtual router.</p> <p>The IP address owner will always become master when available. Preempt mode cannot be disabled on the owner virtual router.</p> <p>The default value for preempt mode is enabled.</p> |
| <b>Default</b>     | preempt                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## priority

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>priority</b> <i>base-priority</i><br><b>no priority</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>router>if>vrrp<br>config>router>if>ipv6>vrrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | <p>This command configures the base router priority for the virtual router instance used in the master election process.</p> <p>The priority is the most important parameter set on a non-owner virtual router instance. The priority defines a virtual router's selection order in the master election process. Together, the priority value and the <b>preempt</b> mode allow the virtual router with the best priority to become the master virtual router.</p> <p>The <i>base-priority</i> is used to derive the in-use priority of the virtual router instance as modified by any optional VRRP priority control policy. VRRP priority control policies can be used to either override or adjust the base priority value depending on events or conditions within the chassis.</p> <p>The <b>priority</b> command is only available in the non-owner <b>vrrp</b> nodal context. The priority of <b>owner</b> virtual router instances is permanently set to 255 and cannot be changed.</p> <p>For non-owner virtual router instances, the default base priority value is 100.</p> <p>The <b>no</b> form of the command reverts to the default value.</p> |

---

|                   |                                                                                                                                                                                                                                              |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | priority 100                                                                                                                                                                                                                                 |
| <b>Parameters</b> | <i>base-priority</i> — The base priority used by the virtual router instance expressed as a decimal integer. If no VRRP priority control policy is defined, the <i>base-priority</i> is the in-use priority for the virtual router instance. |
| <b>Values</b>     | 1 to 254                                                                                                                                                                                                                                     |

## ping-reply

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] ping-reply</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>router>if>vrrp<br>config>router>if>ipv6>vrrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command enables the non-owner master to reply to ICMP echo requests directed at the virtual router instances IP addresses.</p> <p>Non-owner virtual router instances are limited by the VRRP specifications to responding to ARP requests destined to the virtual router IP addresses and routing IP packets not addressed to the virtual router IP addresses. Many network administrators find this limitation frustrating when troubleshooting VRRP connectivity issues.</p> <p>SR OS allows this access limitation to be selectively lifted for certain applications. Ping, Telnet and SSH can be individually enabled or disabled on a per-virtual-router-instance basis.</p> <p>The <b>ping-reply</b> command enables the non-owner master to reply to ICMP echo requests directed at the virtual router instances IP addresses. The Ping request can be received on any routed interface. Ping must not have been disabled at the management security level (either on the parental IP interface or based on the Ping source host address).</p> <p>When <b>ping-reply</b> is not enabled, ICMP echo requests to non-owner master virtual IP addresses are silently discarded.</p> <p>Non-owner backup virtual routers never respond to ICMP echo requests regardless of the <b>ping-reply</b> setting.</p> <p>The <b>ping-reply</b> command is only available in non-owner <b>vrrp</b> nodal context.</p> <p>By default, ICMP echo requests to the virtual router instance IP addresses are silently discarded.</p> <p>The <b>no</b> form of the command configures discarding all ICMP echo request messages destined to the non-owner virtual router instance IP addresses.</p> |
| <b>Default</b>     | no ping-reply — ICMP echo requests to the virtual router instance IP addresses are discarded.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

---

## shutdown

|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>        | <b>[no] shutdown</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>       | config>router>if>vrrp<br>config>router>if>ipv6>vrrp<br>config>vrrp>policy                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>   | <p>This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.</p> <p>The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.</p> <p>The <b>no</b> form of this command administratively enables an entity.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Default</b>       | no shutdown                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Special Cases</b> | <p><b>Non-Owner Virtual Router</b> — Non-owner virtual router instances can be administratively shutdown. This allows the termination of VRRP participation in the virtual router and stops all routing and other access capabilities with regards to the virtual router IP addresses. Shutting down the virtual router instance provides a mechanism to maintain the virtual routers without causing false backup/master state changes.</p> <p>If the <b>shutdown</b> command is executed, no VRRP advertisement messages are generated and all received VRRP advertisement messages are silently discarded with no processing.</p> <p>By default, virtual router instances are created in the <b>no shutdown</b> state.</p> <p>Whenever the administrative state of a virtual router instance transitions, a log message is generated.</p> <p>Whenever the operational state of a virtual router instance transitions, a log message is generated.</p> <p><b>Owner Virtual Router</b> — An owner virtual router context does not have a <b>shutdown</b> command. To administratively disable an owner virtual router instance, use the <b>shutdown</b> command within the parent IP interface node which administratively downs the IP interface.</p> |

## ssh-reply

|                    |                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] ssh-reply</b>                                                                                                                                             |
| <b>Context</b>     | config>router>if>vrrp                                                                                                                                             |
| <b>Description</b> | This command enables the non-owner master to reply to SSH requests directed at the virtual router instance IP addresses. This command is only applicable to IPv4. |



Non-owner virtual router instances are limited by the VRRP specifications to responding to ARP requests destined to the virtual router IP addresses and routing IP packets not addressed to the virtual router IP addresses.

This limitation can be disregarded for certain applications. Ping, Telnet and SSH can be individually enabled or disabled on a per-virtual-router-instance basis.

The **ssh-reply** command enables the non-owner master to reply to SSH requests directed at the virtual router instances IP addresses. The SSH request can be received on any routed interface. SSH must not have been disabled at the management security level (either on the parental IP interface or based on the SSH source host address). Correct login and CLI command authentication is still enforced.

When **ssh-reply** is not enabled, SSH requests to non-owner master virtual IP addresses are silently discarded.

Non-owner backup virtual routers never respond to SSH requests regardless of the **ssh-reply** setting.

The **ssh-reply** command is only available in non-owner **vrrp** nodal context.

By default, SSH requests to the virtual router instance IP addresses are silently discarded.

The **no** form of the command discards all SSH request messages destined to the non-owner virtual router instance IP addresses.

**Default** no ssh-reply — SSH requests to the virtual router instance IP addresses are discarded.

## standby-forwarding

|                    |                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] standby-forwarding</b>                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>router>if>vrrp<br>config>router>if>ipv6>vrrp                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | This command specifies whether this VRRP instance allows forwarding packets to a standby router. When disabled, a standby router should not forward traffic sent to virtual router's MAC address. However, the standby router should forward traffic sent to the standby router's real MAC address. When enabled, a standby router should forward all traffic. |
| <b>Default</b>     | no standby-forwarding                                                                                                                                                                                                                                                                                                                                          |

## telnet-reply

|                |                                                     |
|----------------|-----------------------------------------------------|
| <b>Syntax</b>  | <b>[no] telnet-reply</b>                            |
| <b>Context</b> | config>router>if>vrrp<br>config>router>if>ipv6>vrrp |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command enables the non-owner master to reply to TCP port 23 Telnet requests directed at the virtual router instances' IP addresses.</p> <p>Non-owner virtual router instances are limited by the VRRP specifications to responding to ARP requests destined to the virtual router IP addresses and routing IP packets not addressed to the virtual router IP addresses. Many network administrators find this limitation frustrating when troubleshooting VRRP connectivity issues.</p> <p>This limitation can be disregarded for certain applications. Ping, SSH and Telnet can each be individually enabled or disabled on a per-virtual-router-instance basis.</p> <p>The <b>telnet-reply</b> command enables the non-owner master to reply to Telnet requests directed at the virtual router instances' IP addresses. The Telnet request can be received on any routed interface. Telnet must not have been disabled at the management security level (either on the parental IP interface or based on the Telnet source host address). Correct login and CLI command authentication is still enforced.</p> <p>When <b>telnet-reply</b> is not enabled, Telnet requests to non-owner master virtual IP addresses are silently discarded.</p> <p>Non-owner backup virtual routers never respond to Telnet requests regardless of the <b>telnet-reply</b> setting.</p> <p>The <b>telnet-reply</b> command is only available in non-owner <b>vrrp</b> nodal context.</p> <p>By default, Telnet requests to the virtual router instance IP addresses will be silently discarded.</p> <p>The <b>no</b> form of the command configures discarding all Telnet request messages destined to the non-owner virtual router instance IP addresses.</p> |
| <b>Default</b>     | no telnet-reply — Telnet requests to the virtual router instance IP addresses are discarded.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## traceroute-reply

|                    |                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] traceroute-reply</b>                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>router>if>vrrp<br>config>router>if>ipv6>vrrp                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command is valid only if the VRRP virtual router instance associated with this entry is a non-owner.</p> <p>When this command is enabled, a non-owner master can reply to traceroute requests directed to the virtual router instance IP addresses.</p> <p>A non-owner backup virtual router never responds to such traceroute requests regardless of the <b>trace-route-reply</b> status.</p> |
| <b>Default</b>     | no traceroute-reply                                                                                                                                                                                                                                                                                                                                                                                    |

## vrrp

|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>        | <b>vrrp</b> <i>vrid</i> [ <b>owner</b> ] [ <b>passive</b> ]<br><b>no vrrp</b> <i>vrid</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>       | config>router>interface<br>config>router>if>ipv6                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>   | <p>This command creates the context to configure a VRRP virtual router instance. A virtual router is defined by its virtual router identifier (VRID) and a set of IP addresses.</p> <p>The optional <b>owner</b> keyword indicates that the <b>owner</b> controls the IP address of the virtual router and is responsible for forwarding packets sent to this IP address. The <b>owner</b> assumes the role of the master virtual router.</p> <p>All other virtual router instances participating in this message domain must have the same <i>vrid</i> configured and cannot be configured as <b>owner</b>. Once created, the <b>owner</b> keyword is optional when entering the <i>vrid</i> for configuration purposes.</p> <p>A <i>vrid</i> is internally associated with the IP interface. This allows the <i>vrid</i> to be used on multiple IP interfaces while representing different virtual router instances.</p> <p>For IPv4, up to four VRRP VRID nodes can be configured on a router interface. Each virtual router instance can manage up to 16 backup IP addresses. For IPv6, only one VRID can be configured on a router interface.</p> <p>The optional <b>passive</b> keyword indicates that a <i>vrid</i> can be configured as <b>passive</b>, in which case, the VRRP advertisement messages are suppressed on transmission and reception, and all routers configured with the same <i>vrid</i> become master. Passive VRIDs can exceed the limit of four VRRP VRID nodes on a router interface.</p> <p>The <b>no</b> form of the command removes the specified <i>vrid</i> from the IP interface. This terminates VRRP participation and deletes all references to the <i>vrid</i> in conjunction with the IP interface. The <i>vrid</i> does not need to be shut down to remove the virtual router instance.</p> |
| <b>Default</b>       | no vrrp — No VRRP virtual router instance is associated with the IP interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Special Cases</b> | <p><b>Virtual Router Instance Owner IP Address Conditions</b> — The virtual router instance <b>owner</b> can be created prior to assigning the parent IP interface primary or secondary IP addresses. In this case, the virtual router instance is not associated with an IP address. The operational state of the virtual router instance is down.</p> <p><b>VRRP Owner Command Exclusions</b> — By specifying the VRRP <i>vrid</i> as <b>owner</b>, the following commands are no longer available:</p> <ul style="list-style-type: none"> <li>• <b>vrrp priority</b> — The virtual router instance <b>owner</b> is hard-coded with a <b>priority</b> value of 255 and cannot be changed.</li> <li>• <b>vrrp master-int-inherit</b> — Owner virtual router instances do not accept VRRP advertisement messages; the advertisement interval field is not evaluated and cannot be inherited.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

- **ping-reply, telnet-reply and ssh-reply** — The **owner** virtual router instance always allows Ping, Telnet and SSH if the management and security parameters are configured to accept them on the parent IP interface.
- **vrrp shutdown** — The **owner** virtual router instance cannot be shut down on the **vrrp** node. If this was allowed, VRRP messages would not be sent, but the parent IP interface address would continue to respond to ARPs and forward IP packets. Another virtual router instance may detect the missing master due to the termination of VRRP advertisement messages and become master. This would result in two routers responding to ARP requests for the same IP addresses. To shut down the **owner** virtual router instance, use the **shutdown** command in the parent IP interface context. This will prevent VRRP participation, IP ARP reply and IP forwarding. To continue parent IP interface ARP reply and forwarding without VRRP participation, remove the **vrrp vrid** instance.
- **traceroute-reply**

**VRRP Passive Command Exclusions** — By specifying the VRRP *vrid* as **passive**, the following commands related to the master election and processing of VRRP advertisement messages are no longer available:

- **vrrp priority**
- **policy**
- **preempt**
- **master-int-inherit**
- **standby-forwarding**
- **int-delay**
- **message-interval**
- **authentication-key**
- **bfd-enable**

**Parameters**     *vrid* — The virtual router ID for the IP interface expressed as a decimal integer.

**Values**            1 to 255

**owner** — Identifies this virtual router instance as owning the virtual router IP addresses. If the **owner** keyword is not specified at the time of *vrid* creation, the **vrrp backup** commands must be specified to define the virtual router IP addresses. The **owner** keyword is not required when entering the *vrid* for editing purposes. When created as **owner**, a *vrid* on an IP interface cannot have the **owner** parameter removed. The *vrid* must be deleted, and then recreated without the **owner** keyword, to remove ownership.

**passive** — Identifies this virtual router instance as **passive**, therefore owning the virtual router IP addresses. A **passive vrid** does not send or receive VRRP advertisement messages and is always in either the **master** state (if the interface is operationally up), or the **init** state (if the interface is operationally down). The **passive** keyword is not required when entering the *vrid* for editing purposes. When a *vrid* on an IP interface is created as **passive**, the parameter cannot be removed from the *vrid*. The *vrid* must be deleted, and then recreated without the **passive** keyword, to remove the parameter.

### 3.9.2.2 Priority Policy Commands

#### delta-in-use-limit

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>delta-in-use-limit</b> <i>in-use-priority-limit</i><br><b>no delta-in-use-limit</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>vrrp>policy                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <p>This command sets a lower limit on the virtual router in-use priority that can be derived from the delta priority control events.</p> <p>Each <i>vrrp-priority-id</i> places limits on the delta priority control events to define the in-use priority of the virtual router instance. Setting this limit prevents the sum of the delta priority events from lowering the in-use priority value of the associated virtual router instances below the configured value.</p> <p>The limit has no effect on explicit priority control events. Explicit priority control events are controlled by setting the in-use priority to any value between 1 and 254.</p> <p>Only non-owner virtual router instances can be associated with VRRP priority control policies and their priority control events.</p> <p>Once the total sum of all delta events is calculated and subtracted from the base <b>priority</b> of the virtual router instance, the result is compared to the <b>delta-in-use-limit</b> value. If the result is less than the limit, the <b>delta-in-use-limit</b> value is used as the virtual router in-use priority value. If an explicit priority control event overrides the delta priority control events, the <b>delta-in-use-limit</b> has no effect.</p> <p>Setting the limit to a higher value than the default of 1 limits the effect of the delta priority control events on the virtual router instance base <b>priority</b> value. This allows for multiple priority control events while minimizing the overall effect on the in-use priority.</p> <p>Changing the <i>in-use-priority-limit</i> causes an immediate re-evaluation of the in-use priority values for all virtual router instances associated with this <i>vrrp-policy-id</i> based on the current sum of all active delta control policy events.</p> <p>The <b>no</b> form of the command reverts to the default value.</p> |
| <b>Default</b>     | delta-in-use-limit 1 — The lower limit of 1 for the in-use priority, as modified, by delta priority control events.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Parameters</b>  | <i>in-use-priority-limit</i> — The lower limit of the in-use priority base, as modified by priority control policies. The <i>in-use-priority-limit</i> has the same range as the non-owner virtual router instance base-priority parameter. If the result of the total delta priority control events minus the virtual router instances base-priority, is less than the <i>in-use-priority-limit</i> , the <i>in-use-priority-limit</i> value is used as the virtual router instances in-use priority value.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

Setting the *in-use-priority-limit* to a value equal to or larger than the virtual router instance *base-priority* prevents the delta priority control events from having any effect on the virtual router instance in-use priority value.

**Values** 1 to 254

## description

|                    |                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>description</b> <i>string</i><br><b>no description</b>                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>vrrp>policy                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command creates a text description stored in the configuration file for a configuration context.</p> <p>The <b>description</b> command associates a text string with a configuration context to help identify the content in the configuration file.</p> <p>The <b>no</b> form of the command removes the string from the configuration.</p> |
| <b>Default</b>     | n/a                                                                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <i>string</i> — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes.                                                                    |

## policy

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>policy</b> <i>policy-id</i> [ <b>context</b> <i>service-id</i> ]<br><b>no policy</b> <i>policy-id</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>vrrp                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command creates the context to configure a VRRP priority control policy which is used to control the VRRP in-use priority based on priority control events. It is a parental node for the various VRRP priority control policy commands that define the policy parameters and priority event conditions.</p> <p>The virtual router instance <b>priority</b> command defines the initial or base value to be used by non-owner virtual routers. This value can be modified by assigning a VRRP priority control policy to the virtual router instance. The VRRP priority control policy can override or diminish the base priority setting to establish the actual in-use priority of the virtual router instance.</p> <p>The <b>policy</b> <i>policy-id</i> command must be created first, before it can be associated with a virtual router instance.</p> |

Because VRRP priority control policies define conditions and events that must be maintained, they can be resource intensive. The number of policies is limited to 1000.

The *policy-id* do not have to be consecutive integers. The range of available policy identifiers is from 1 to 9999.

The **no** form of the command deletes the specific *policy-id* from the system. The *policy-id* must be removed first from all virtual router instances before the **no policy** command can be issued. If the *policy-id* is associated with a virtual router instance, the command will fail.

|                   |                                                                                                                                                                                                                                 |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>vrrp-policy-id</i> — The VRRP priority control ID expressed as a decimal integer that uniquely identifies this policy from any other VRRP priority control policy defined on the system. Up to 1000 policies can be defined. |
|                   | <b>Values</b> 1 to 9999                                                                                                                                                                                                         |
|                   | <b>context service-id</b> — Specifies the service ID to which this policy applies. A value of zero (0) means that this policy does not apply to a service but applies to the base router instance.                              |
|                   | <b>Values</b> 1 to 2147483647                                                                                                                                                                                                   |

## priority-event

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] priority-event</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>vrrp>policy>priority-event                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command creates the context to configure VRRP priority control events used to define criteria to modify the VRRP in-use priority.</p> <p>A priority control event specifies an object to monitor and the effect on the in-use priority level for an associated virtual router instance.</p> <p>Up to 32 priority control events can be configured within the <b>priority-event</b> node.</p> <p>The <b>no</b> form of the command clears any configured priority events.</p> |
| <b>Default</b>     | n/a                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

### 3.9.2.3 Priority Policy Event Commands

## hold-clear

|               |                                                   |
|---------------|---------------------------------------------------|
| <b>Syntax</b> | <b>hold-clear seconds</b><br><b>no hold-clear</b> |
|---------------|---------------------------------------------------|

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | <pre>config&gt;vrrp&gt;policy&gt;priority-event&gt;host-unreachable config&gt;vrrp&gt;policy&gt;priority-event&gt;lag-port-down config&gt;vrrp&gt;policy&gt;priority-event&gt;mc-ipsec-non-forwarding config&gt;vrrp&gt;policy&gt;priority-event&gt;port-down config&gt;vrrp&gt;policy&gt;priority-event&gt;route-unknown</pre>                                                                                                                                                                    |
| <b>Description</b> | <p>This command configures the hold clear time for the event. The <i>seconds</i> parameter specifies the hold-clear time, the amount of time in seconds by which the effect of a cleared event on the associated virtual router instance is delayed.</p> <p>The hold-clear time is used to prevent black hole conditions when a virtual router instance advertises itself as a master before other conditions associated with the cleared event have had a chance to enter a forwarding state.</p> |
| <b>Default</b>     | no hold-clear                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Parameters</b>  | <p><i>seconds</i> — Specifies the amount of time in seconds by which the effect of a cleared event on the associated virtual router instance is delayed.</p> <p><b>Values</b>      0 to 86400</p>                                                                                                                                                                                                                                                                                                  |

## hold-set

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <pre><b>hold-set</b> <i>seconds</i> <b>no hold-set</b></pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | <pre>config&gt;vrrp&gt;policy&gt;priority-event&gt;host-unreachable config&gt;vrrp&gt;policy&gt;priority-event&gt;lag-port-down config&gt;vrrp&gt;policy&gt;priority-event&gt;mc-ipsec-non-forwarding config&gt;vrrp&gt;policy&gt;priority-event&gt;port-down config&gt;vrrp&gt;policy&gt;priority-event&gt;route-unknown</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command specifies the amount of time that must pass before the set state for a VRRP priority control event can transition to the cleared state to dampen flapping events. A flapping event continually transitions between clear and set.</p> <p>The <b>hold-set</b> command is used to dampen the effect of a flapping event. The <b>hold-set</b> value is loaded into a hold-set timer that prevents a set event from transitioning to the cleared state until it expires.</p> <p>Each time an event transitions between cleared and set, the timer is loaded and begins a countdown to zero. When the timer reaches zero, the event is allowed to enter the cleared state. Entering the cleared state is dependent on the object controlling the event, conforming to the requirements defined in the event itself. It is possible, on some event types, to have another set action reload the hold-set timer. This extends the amount of time that must expire before entering the cleared state.</p> |



Once the hold-set timer expires and the event meets the cleared state requirements or is set to a lower threshold, the current set effect on the virtual router instances in-use priority can be removed. As with **lag-port-down** events, this may be a decrease in the set effect if the *clearing* amounts to a lower set threshold.

The **hold-set** command can be executed at anytime. If the hold-set timer value is configured larger than the new *seconds* setting, the timer is loaded with the new **hold-set** value.

The **no** form of the command reverts the default value.

**Default** 0 — The hold-set timer is disabled so event transitions are processed immediately.

**Parameters** *seconds* — The number of seconds that the hold-set timer waits after an event enters a set state or enters a higher threshold set state, depending on the event type.  
The value of 0 disables the hold-set timer, preventing any delay in processing lower set thresholds or cleared events.

**Values** 0 to 86400

priority

**Syntax** **priority** *priority-level* [{**delta** | **explicit**}]  
**no priority**

**Context** config>vrrp>policy>priority-event>host-unreachable  
config>vrrp>policy>priority-event>lag-port-down  
config>vrrp>policy>priority-event>mc-ipsec-non-forwarding  
config>vrrp>policy>priority-event>port-down  
config>vrrp>policy>priority-event>route-unknown

**Description** This command controls the effect the set event has on the virtual router instance in-use priority.

When the event is set, the *priority-level* is either subtracted from the base priority of each virtual router instance or it defines the explicit in-use priority value of the virtual router instance depending on whether the **delta** or **explicit** keywords are specified.

Multiple set events in the same policy have interaction constraints:

- If any set events have an explicit **priority** value, all the delta **priority** values are ignored.
- The set event with the lowest explicit **priority** value defines the in-use priority that are used by all virtual router instances associated with the policy.
- If no set events have an explicit **priority** value, all the set events delta **priority** values are added and subtracted from the base priority value defined on each virtual router instance associated with the policy.
- If the delta priorities sum exceeds the **delta-in-use-limit** parameter, then the **delta-in-use-limit** parameter is used as the value subtracted from the base priority value defined on each virtual router instance associated with the policy.

If the **priority** command is not configured on the priority event, the *priority-value* defaults to 0 and the qualifier keyword defaults to **delta**, therefore, there is no impact on the in-use priority.

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

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | 0 delta — The set event will subtract 0 from the base priority (no effect).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b> | <i>priority-level</i> — The priority level adjustment value expressed as a decimal integer.<br><b>Values</b> 0 to 254<br><b>delta   explicit</b> — Configures what effect the <i>priority-level</i> will have on the base priority value.<br>When <b>delta</b> is specified, the <i>priority-level</i> value is subtracted from the associated virtual router instance's base priority when the event is set and no explicit events are set. The sum of the priority event <i>priority-level</i> values on all set delta priority events are subtracted from the virtual router base priority to derive the virtual router instance in-use priority value. If the <b>delta</b> priority event is cleared, the <i>priority-level</i> is no longer used in the in-use priority calculation.<br>When <b>explicit</b> is specified, the <i>priority-level</i> value is used to override the base priority of the virtual router instance if the priority event is set and no other <b>explicit</b> priority event is set with a lower <i>priority-level</i> . The set <b>explicit</b> priority value with the lowest <i>priority-level</i> determines the actual in-use protocol value for all virtual router instances associated with the policy.<br><b>Default</b> delta<br><b>Values</b> delta, explicit |

weight-down

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>weight-down</b> <i>lag-ports-down-weight</i><br><b>no weight-down</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>vrrp>policy>priority-event>lag-port-down                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | This command creates a context to configure an event set threshold within a lag-port-down priority control event. The weight-down command defines a sub-node within the lag-port-down event and is uniquely identified with the lag-ports-down-weight parameter. Each weight-down node within the same lag-port-down event node must have a unique lag-ports-down-weight value. Each weight-down node has its own priority command that takes effect whenever that node represents the current threshold. A single LAG can use either weight-threshold or port threshold. The command is required for correct operation on mixed port-speed LAGs and can be used for non-mixed port-speed LAGs as well.<br><br>The total number of sub-nodes (uniquely identified by the lag-ports-down-weight parameter) allowed in the system is 2048. |

A **weight-down** node is not required for each possible number of ports that could be down. The active threshold is always the closest lower threshold.

The **no** form of the command deletes the event set threshold. The threshold may be removed at any time. If the removed threshold is the current active threshold, the event set thresholds must be re-evaluated after removal.

**Default** no weight-down

**Parameters** *lag-ports-down-weight* — The total weight of LAG ports down to create a set event threshold. This is the active threshold when the weight of down ports in the LAG equals or exceeds *lag-ports-down-weight*, but does not equal or exceed the next highest configured *lag-ports-down-weight*.

**Values** 1 to 64

## mc-ipsec-non-forwarding

**Syntax** [no] mc-ipsec-non-forwarding *tunnel-grp-id*

**Context** config>vrrp>policy>priority-event

**Description** This command configures an instance of a multi-chassis IPsec tunnel-group Priority Event used to override the base priority value of a VRRP virtual router instance depending on the operational state of the event.

**Default** n/a

**Parameters** *tunnel-grp-id* — Identifies the multi-chassis IPsec tunnel group whose non-forwarding state is monitored by this priority control event.

### 3.9.2.4 Priority Policy Port Down Event Commands

## port-down

**Syntax** [no] port-down *port-id*

**Context** config>vrrp>policy>priority-event

**Description** This command configures a port down priority control event that monitors the operational state of a port or SONET/SDH channel. When the port or channel enters the operational down state, the event is considered set. When the port or channel enters the operational up state, the event is considered cleared.

Multiple unique **port-down** event nodes can be configured within the **priority-event** context up to the overall limit of 32 events. Up to 32 events can be defined in any combination of types.

The **port-down** command can reference an arbitrary port or channel. The port or channel does not need to be preprovisioned or populated within the system. The operational state of the **port-down** event is set as follows:

- Set – non-provisioned
- Set – not populated
- Set – down
- Cleared – up

When the port or channel is provisioned, populated, or enters the operationally up or down state, the event operational state is updated appropriately.

When the event enters the operationally down, non-provisioned, or non-populated state, the event is considered to be set. When an event transitions from clear to set, the set is processed immediately and must be reflected in the associated virtual router instances in-use priority value. As the event transitions from cleared to set, a hold-set timer is loaded with the value configured by the events **hold-set** command. This timer prevents the event from clearing until it expires, damping the effect of event flapping. If the event clears and becomes set again before the hold-set timer expires, the timer is reset to the **hold-set** value, extending the time before another clear can take effect.

When the event enters the operationally up state, the event is considered to be cleared. Once the events **hold-set** expires, the effects of the events **priority** value are immediately removed from the in-use priority of all associated virtual router instances.

The actual effect on the virtual router instance in-use priority value depends on the defined event priority and its delta or explicit nature.

The **no** form of the command deletes the specific port or channel monitoring event. The event may be removed at anytime. When the event is removed, the in-use priority of all associated virtual router instances will be re-evaluated. The events **hold-set** timer has no effect on the removal procedure.

**Default** no port-down — No port down priority control events are defined.

**Parameters** *port-id* — The port ID of the port monitored by the VRRP priority control event.

The *port-id* can only be monitored by a single event in this policy. The port can be monitored by multiple VRRP priority control policies. A port and a specific channel on the port are considered to be separate entities. A port and a channel on the port can be monitored by separate events in the same policy.

**Values** The following values apply to the 7750 SR:

|                                   |                                  |                    |  |
|-----------------------------------|----------------------------------|--------------------|--|
| port-id                           | <i>slot/mda/port[.channel]</i>   |                    |  |
| eth-sat-id                        | <i>esat-id/slot/port</i>         |                    |  |
|                                   | <i>esat</i>                      | keyword            |  |
|                                   | <i>id</i>                        | 1 to 20            |  |
| pxc-id                            | <i>pxc-id.sub-port</i>           |                    |  |
|                                   | <i>pxc</i>                       | keyword            |  |
|                                   | <i>id</i>                        | 1 to 64            |  |
|                                   | <i>sub-port</i>                  | a, b               |  |
| aps-id                            | <i>aps-group-id[.channel]</i>    |                    |  |
|                                   | <i>aps</i>                       | keyword            |  |
|                                   | <i>group-id</i>                  | 1 to 64            |  |
| bundle-type-slot/mda.<bundle-num> |                                  |                    |  |
|                                   | <i>bundle</i>                    | keyword            |  |
|                                   | <i>type</i>                      | ima, ppp           |  |
|                                   | <i>bundle-num</i>                | 1 to 256           |  |
| ccag-id                           | <i>ccag-id. path-id[cc-type]</i> |                    |  |
|                                   | <i>ccag</i>                      | keyword            |  |
|                                   | <i>id</i>                        | 1 to 8             |  |
|                                   | <i>path-id</i>                   | a, b               |  |
|                                   | <i>cc-type</i>                   | .sap-net, .net-sap |  |

**Values** The following values apply to the 7450 ESS:

|            |                                  |                    |  |
|------------|----------------------------------|--------------------|--|
| port-id    | <i>slot/mda/</i>                 |                    |  |
|            | <i>port[.channel]</i>            |                    |  |
| eth-sat-id | <i>esat-id/slot/port</i>         |                    |  |
|            | <i>esat</i>                      | keyword            |  |
|            | <i>id</i>                        | 1 to 20            |  |
| pxc-id     | <i>pxc-id.sub-port</i>           |                    |  |
|            | <i>pxc</i>                       | keyword            |  |
|            | <i>id</i>                        | 1 to 64            |  |
|            | <i>sub-port</i>                  | a, b               |  |
| ccag-id    | <i>ccag-id. path-id[cc-type]</i> |                    |  |
|            | <i>ccag</i>                      | keyword            |  |
|            | <i>id</i>                        | 1 to 8             |  |
|            | <i>path-id</i>                   | a, b               |  |
|            | <i>cc-type</i>                   | .sap-net, .net-sap |  |

The POS channel on the port monitored by the VRRP priority control event. The *port-id.channel-id* can only be monitored by a single event in this policy. The channel can be monitored by multiple VRRP priority control policies. A port and a specific channel on the port are considered to be separate entities. A port and a channel on the port can be monitored by separate events in the same policy.

If the port is provisioned, but the *channel* does not exist or the port has not been populated, the appropriate event operational state is Set – non-populated.

If the port is not provisioned, the event operational state is Set – non-provisioned.

If the POS interface is configured as a clear-channel, the *channel-id* is 1 and the channel bandwidth is the full bandwidth of the port.

### 3.9.2.5 Priority Policy LAG Events Commands

#### lag-port-down

|                    |                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] lag-port-down</b> <i>lag-id</i>                                                                                                                        |
| <b>Context</b>     | config>vrrp>policy>priority-event                                                                                                                              |
| <b>Description</b> | This command creates the context to configure Link Aggregation Group (LAG) priority control events that monitor the operational state of the links in the LAG. |

The **lag-port-down** command configures a priority control event. The event monitors the operational state of each port in the specified LAG. When one or more of the ports enter the operational down state, the event is considered to be set. When all the ports enter the operational up state, the event is considered to be clear. As ports enter the operational up state, any previous set threshold that represents more down ports is considered cleared, while the event is considered to be set.

Multiple unique **lag-port-down** event nodes can be configured within the **priority-event** node up to the maximum of 32 events.

The **lag-port-down** command can reference an arbitrary LAG. The *lag-id* does have to already exist within the system. The operational state of the **lag-port-down** event will indicate:

- Set – non-existent
- Set – one port down
- Set – two ports down
- Set – three ports down
- Set – four ports down
- Set – five ports down

- Set – six ports down
- Set – seven ports down
- Set – eight ports down
- Cleared – all ports up

When the *lag-id* is created, or a port in *lag-id* becomes operationally up or down, the event operational state must be updated appropriately.

When one or more of the LAG composite ports enters the operationally down state or the *lag-id* is deleted or does not exist, the event is considered to be set. When an event transitions from clear to set, the set is processed immediately and must be reflected in the associated virtual router instances in-use priority value. As the event transitions from clear to set, a hold-set timer is loaded with the value configured by the events **hold-set** command. This timer prevents the event from clearing until it expires, damping the effect of event flapping. If the event clears and becomes set again before the hold-set timer expires, the timer is reset to the **hold-set** value, extending the time before another clear can take effect.

The **lag-port-down** event is considered to have a tiered event set state. While the priority impact per number of ports down is totally configurable, as more ports go down, the effect on the associated virtual router instances in-use priority is expected to increase (lowering the priority). When each configured threshold is crossed, any higher thresholds are considered further event sets and are processed immediately with the hold-set timer reset to the configured value of the **hold-set** command. As the thresholds are crossed in the opposite direction (fewer ports down then previously), the priority effect of the event is not processed until the hold-set timer expires. If the number of ports down threshold again increases before the hold-set timer expires, the timer is only reset to the **hold-set** value if the number of ports down is equal to or greater than the threshold that set the timer.

The event contains **number-down** nodes that define the priority delta or explicit value to be used based on the number of LAG composite ports that are in the operationally down state. These nodes represent the event set thresholds. Not all port down thresholds must be configured. As the number of down ports increase, the **number-down ports-down** node that expresses a value equal to or less than the number of down ports describes the delta or explicit priority value to be applied.

The **no** form of the command deletes the specific LAG monitoring event. The event can be removed at anytime. When the event is removed, the in-use priority of all associated virtual router instances must be reevaluated. The events **hold-set** timer has no effect on the removal procedure.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no lag-port-down — No LAG priority control events are created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b> | <p><i>lag-id</i> — The LAG ID that the specific event is to monitor expressed as a decimal integer. The <i>lag-id</i> can only be monitored by a single event in this policy. The LAG may be monitored by multiple VRRP priority control policies. A port within the LAG and the LAG ID itself are considered to be separate entities. A composite port may be monitored with the <b>port-down</b> event while the <i>lag-id</i> the port is in is monitored by a <b>lag-port-down</b> event in the same policy.</p> <p><b>Values</b> 1 to 800 (apply to the 7750 SR and 7950 XRS)</p> |

1 to 200 (apply to the 7450 ESS)

number-down

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] number-down</b> <i>number-of-lag-ports-down</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>vrrp>policy>priority-event>lag-port-down                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>This command creates a context to configure an event set threshold within a lag-port-down priority control event.</p> <p>The <b>number-down</b> command defines a sub-node within the <b>lag-port-down</b> event and is uniquely identified with the <i>number-of-lag-ports-down</i> parameter. Each <b>number-down</b> node within the same <b>lag-port-down</b> event node must have a unique <i>number-of-lag-ports-down</i> value. Each <b>number-down</b> node has its own <b>priority</b> command that takes effect whenever that node represents the current threshold.</p> <p>The total number of sub-nodes (uniquely identified by the <i>number-of-lag-ports-down</i> parameter) allowed in a single <b>lag-port-down</b> event is equal to the total number of possible physical ports allowed in a LAG.</p> <p>A <b>number-down</b> node is not required for each possible number of ports that could be down. The active threshold is always the closest lower threshold. When the number of ports down equals a given threshold, that is the active threshold.</p> <p>The <b>no</b> form of the command deletes the event set threshold. The threshold may be removed at any time. If the removed threshold is the current active threshold, the event set thresholds must be re-evaluated after removal.</p> |
| <b>Default</b>     | no number-down — No threshold for the LAG priority event is created.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Parameters</b>  | <p><i>number-of-lag-ports-down</i> — The number of LAG ports down to create a set event threshold. This is the active threshold when the number of down ports in the LAG equals or exceeds <i>number-of-lag-ports-down</i>, but does not equal or exceed the next highest configured <i>number-of-lag-ports-down</i>.</p> <p><b>Values</b>      1 to 64 (applies to 64-link LAG)<br/>                 1 to 32 (applies to other LAGs)</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

3.9.2.6 Priority Policy Host Unreachable Event Commands

drop-count

|               |                                                                       |
|---------------|-----------------------------------------------------------------------|
| <b>Syntax</b> | <b>drop-count</b> <i>consecutive-failures</i><br><b>no drop-count</b> |
|---------------|-----------------------------------------------------------------------|



---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>vrrp vrrp-policy-id>priority-event>host-unreachable                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command configures the number of consecutively sent ICMP echo request messages that must fail before the host unreachable priority control event is set.</p> <p>The <b>drop-count</b> command is used to define the number of consecutive message send attempts that must fail for the <b>host-unreachable</b> priority event to enter the set state. Each unsuccessful attempt increments the event's consecutive message drop counter. With each successful attempt, the event's consecutive message drop counter resets to zero.</p> <p>If the event's consecutive message drop counter reaches the <b>drop-count</b> value, the <b>host-unreachable</b> priority event enters the set state.</p> <p>The event's <b>hold-set</b> value defines how long the event must stay in the set state even when a successful message attempt clears the consecutive drop counter. The event is not cleared until the consecutive drop counter is less than the <b>drop-count</b> value and the <b>hold-set</b> timer has a value of zero (expired).</p> <p>The <b>no</b> form of the command reverts to the default value.</p> |
| <b>Default</b>     | drop-count 3 — 3 consecutive ICMP echo request failures are required before the host unreachable priority control event is set.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Parameters</b>  | <p><i>consecutive-failures</i> — The number of ICMP echo request message attempts that must fail for the event to enter the set state. It also defines the threshold so a lower consecutive number of failures can clear the event state.</p> <p><b>Values</b>      1 to 60</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## host-unreachable

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p>[no] <b>host-unreachable</b> <i>ip-address</i></p> <p>[no] <b>host-unreachable</b> <i>ipv6-address</i></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>vrrp>policy>priority-event                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command creates the context to configure a host unreachable priority control event to monitor the ability to receive ICMP echo reply packets from an IP host address.</p> <p>A host unreachable priority event creates a continuous ICMP echo request (ping) probe to the specified <i>ip-address</i>. If a ping fails, the event is considered to be set. If a ping is successful, the event is considered to be cleared.</p> <p>Multiple unique (different <i>ip-address</i>) <b>host-unreachable</b> event nodes can be configured within the <b>priority-event</b> node to a maximum of 32 events.</p> |

The **host-unreachable** command can reference any valid local or remote IP address. The ability to ARP a local IP address or find a remote IP address within a route prefix in the route table is considered part of the monitoring procedure. The **host-unreachable** priority event operational state tracks ARP or route table entries dynamically appearing and disappearing from the system. The operational state of the **host-unreachable** event are listed in [Table 40](#).

**Table 40 Host Unreachable Operational States**

| Host Unreachable Operational State | Description                                                                                                                           |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| Set – no ARP                       | No ARP address found for <i>ip-addr</i> for <b>drop-count</b> consecutive attempts. Only applies when IP address is considered local. |
| Set – no route                     | No route exists for <i>ip-addr</i> for <b>drop-count</b> consecutive attempts. Only when IP address is considered remote.             |
| Set – host unreachable             | ICMP host unreachable message received for <b>drop-count</b> consecutive attempts.                                                    |
| Set – no reply                     | ICMP echo request timed out for <b>drop-count</b> consecutive attempts.                                                               |
| Set – reply received               | Last ICMP echo request attempt received an echo reply but historically not able to clear the event.                                   |
| Cleared – no ARP                   | No ARP address found for <i>ip-addr</i> - not enough failed attempts to set the event.                                                |
| Cleared – no route                 | No route exists for <i>ip-addr</i> - not enough failed attempts to set the event.                                                     |
| Cleared – host unreachable         | ICMP host unreachable message received - not enough failed attempts to set the event.                                                 |
| Cleared – no reply                 | ICMP echo request timed out - not enough failed attempts to set the event.                                                            |
| Cleared – reply received           | Event is cleared - last ICMP echo request received an echo reply.                                                                     |

Unlike other priority event types, the **host-unreachable** priority event monitors a repetitive task. A historical evaluation is performed on the success rate of receiving ICMP echo reply messages. The operational state takes its cleared and set orientation from the historical success rate. The informational portion of the operational state is derived from the last attempt's result. It is possible for the previous attempt to fail while the operational state is still cleared due to an insufficient number of failures to cause it to become set. It is also possible for the state to be set while the previous attempt was successful.

When an event transitions from clear to set, the set is processed immediately and must be reflected in the associated virtual router instances in-use priority value. As the event transitions from clear to set, a hold-set timer is loaded with the value configured by the events **hold-set** command. This timer prevents the event from clearing until it expires, damping the effect of event flapping. If the event clears and becomes set again before the hold-set timer expires, the timer is reset to the **hold-set** value, extending the time before another clear can take effect.

The hold-set timer be expired and the historical success rate must be met prior to the event operational state becoming cleared.

The **no** form of the command deletes the specific IP host monitoring event. The event may be deleted at anytime. When the event is deleted, the in-use priority of all associated virtual router instances must be reevaluated. The event's **hold-set** timer has no effect on the removal procedure.

**Default** no host-unreachable — No host unreachable priority events are created.

**Parameters** *ip-addr* — The IP address of the host for which the specific event will monitor connectivity. The *ip-addr* can only be monitored by a single event in this policy. The IP address can be monitored by multiple VRRP priority control policies. The IP address can be used in one or multiple **ping** requests. Each VRRP priority control **host-unreachable** and **ping** destined to the same *ip-addr* is uniquely identified on a per message basis. Each session originates a unique identifier value for the ICMP echo request messages it generates. This allows received ICMP echo reply messages to be directed to the appropriate sending application.

**Values** The following values apply to the 7450 ESS:  
ipv4-address: a.b.c.d

**Values** The following values apply to the 7750 SR and 7950 XRS:  
ipv4-address: a.b.c.d  
ipv6-address: x:x:x:x:x:x[-interface]  
x: [0..FFFF]H  
interface: 32 chars maximum, mandatory for link local addresses

The link-local IPv6 address must have an interface name specified. The global IPv6 address must not have an interface name specified.

interval

**Syntax** interval seconds  
no interval

---

|                    |                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>vrrp>priority-event>host-unreachable                                                                                                                                                                                     |
| <b>Description</b> | <p>This command configures the number of seconds between host unreachable priority event ICMP echo request messages directed to the host IP address.</p> <p>The <b>no</b> form of the command reverts to the default value.</p> |
| <b>Default</b>     | interval 1                                                                                                                                                                                                                      |
| <b>Parameters</b>  | <p><i>seconds</i> — The number of seconds between the ICMP echo request messages sent to the host IP address for the host unreachable priority event.</p> <p><b>Values</b> 1 to 60</p>                                          |

## padding-size

|                    |                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>padding-size</b> <i>size</i><br><b>no padding-size</b>                                                                                                      |
| <b>Context</b>     | config>vrrp>priority-event>host-unreachable                                                                                                                    |
| <b>Description</b> | <p>This command allows the operator to increase the size of IP packet by padding the PDU.</p> <p>The <b>no</b> form of the command reverts to the default.</p> |
| <b>Default</b>     | padding-size 0                                                                                                                                                 |
| <b>Parameters</b>  | <p><i>size</i> — Specifies amount of increase to the ICMP PDU.</p> <p><b>Values</b> 0 to 16384</p>                                                             |

## timeout

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>timeout</b> <i>seconds</i><br><b>no timeout</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Context</b>     | config>vrrp vrrp-policy-id>priority-event>host-unreachable                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command defines the time, in seconds, that must pass before considering the far-end IP host unresponsive to an outstanding ICMP echo request message.</p> <p>The <b>timeout</b> value is not directly related to the configured <b>interval</b> parameter. The <b>timeout</b> value may be larger, equal, or smaller, relative to the <b>interval</b> value.</p> <p>If the <b>timeout</b> value is larger than the <b>interval</b> value, multiple ICMP echo request messages may be outstanding. Every ICMP echo request message transmitted to the far end host is tracked individually according to the message identifier and sequence number.</p> <p>With each consecutive attempt to send an ICMP echo request message, the timeout timer is loaded with the <b>timeout</b> value. The timer decrements until:</p> |

- An internal error occurs preventing message sending (request unsuccessful).
- An internal error occurs preventing message reply receiving (request unsuccessful).
- A required route table entry does not exist to reach the IP address (request unsuccessful).
- A required ARP entry does not exist and ARP request timed out (request unsuccessful).
- A valid reply is received (request successful).

It is possible for a required ARP request to succeed or timeout after the message timeout timer expires. In this case, the message request is unsuccessful.

If an ICMP echo reply message is not received prior to the **timeout** period for a given ICMP echo request, that request is considered to be dropped and increments the consecutive message drop counter for the priority event.

If an ICMP echo reply message with the same sequence number as an outstanding ICMP echo request message is received prior to that message timing out, the request is considered successful. The consecutive message drop counter is cleared and the request message no longer is outstanding.

If an ICMP Echo Reply message with a sequence number equal to an ICMP echo request sequence number that had previously timed out is received, that reply is silently discarded while incrementing the priority event reply discard counter.

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

|                   |                                                                                                                                                                                          |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | timeout 1                                                                                                                                                                                |
| <b>Parameters</b> | <i>seconds</i> — The number of seconds before an ICMP echo request message is timed out. Once a message is timed out, a reply with the same identifier and sequence number is discarded. |
| <b>Values</b>     | 1 to 60                                                                                                                                                                                  |

### 3.9.2.7 Priority Policy Route Unknown Event Commands

#### less-specific

|                    |                                                                                                                                                     |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [no] <b>less-specific</b> [allow-default]                                                                                                           |
| <b>Context</b>     | config>vrrp>policy>priority-event>route-unknown                                                                                                     |
| <b>Description</b> | This command allows a CIDR shortest match hit on a route prefix that contains the IP route prefix associated with the route unknown priority event. |

The **less-specific** command modifies the search parameters for the IP route prefix specified in the **route-unknown** priority event. Specifying **less-specific** allows a CIDR shortest match hit on a route prefix that contains the IP route prefix.

The **less-specific** command eases the RTM lookup criteria when searching for the *prefix/mask-length*. When the **route-unknown** priority event sends the prefix to the RTM (as if it was a destination lookup), the result route table prefix (if a result is found) is checked to see if it is an exact match or a less specific match. The **less-specific** command enables a less specific route table prefix to match the configured prefix. When **less-specific** is not specified, a less specific route table prefix fails to match the configured prefix. The **allow-default** optional parameter extends the **less-specific** match to include the default route (0.0.0.0).

The **no** form of the command prevents RTM lookup results that are less specific than the route prefix from matching.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default</b>    | no less-specific — The route unknown priority events requires an exact prefix/mask match.                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b> | <b>allow-default</b> — When the <b>allow-default</b> parameter is specified with the <b>less-specific</b> command, an RTM return of 0.0.0.0 matches the IP prefix. If <b>less-specific</b> is entered without the <b>allow-default</b> parameter, a return of 0.0.0.0 will not match the IP prefix. To disable <b>allow-default</b> , but continue to allow <b>less-specific</b> match operation, only enter the <b>less-specific</b> command (without the <b>allow-default</b> parameter). |

## next-hop

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [no] <b>next-hop</b> <i>ip-address</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Context</b>     | config>vrrp>policy>priority-event>route-unknown                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command adds an allowed next hop IP address to match the IP route prefix for a route-unknown priority control event.</p> <p>If the next-hop IP address does not match one of the defined <i>ip-address</i>, the match is considered unsuccessful and the <b>route-unknown</b> event transitions to the set state.</p> <p>The <b>next-hop</b> command is optional. If no <b>next-hop</b> <i>ip-address</i> commands are configured, the comparison between the RTM prefix return and the <b>route-unknown</b> IP route prefix are not included in the next hop information.</p> <p>When more than one next hop IP addresses are eligible for matching, a <b>next-hop</b> command must be executed for each IP address. Defining the same IP address multiple times has no effect after the first instance.</p> <p>The <b>no</b> form of the command removes the <i>ip-address</i> from the list of acceptable next hops when looking up the <b>route-unknown</b> prefix. If this <i>ip-address</i> is the last next hop defined on the <b>route-unknown</b> event, the returned next hop information is ignored when testing the match criteria. If the <i>ip-address</i> does not exist, the <b>no next-hop</b> command returns a warning error, but continues to execute if part of an <b>exec</b> script.</p> |
| <b>Default</b>     | no next-hop — No next hop IP address for the route unknown priority control event is defined.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

**Parameters** *ip-address* — The IP address for an acceptable next hop IP address for a returned route prefix from the RTM when looking up the **route-unknown** route prefix.

**Values** The following values apply to the 7450 ESS:  
ipv4-address: a.b.c.d

**Values** The following values apply to the 7750 SR and 7950 XRS:

|               |                                                      |
|---------------|------------------------------------------------------|
| ipv4-address: | a.b.c.d                                              |
| ipv6-address: | x:x:x:x:x:x[-interface]                              |
| x:            | [0..FFFF]H                                           |
| interface:    | 32 chars maximum, mandatory for link local addresses |

The link-local IPv6 address must have an interface name specified. The global IPv6 address must not have an interface name specified.

protocol

**Syntax** **protocol** {**bgp** | **bgp-vpn** | **ospf** | **is-is** | **rip** | **static**}  
**no protocol**

**Context** config>vrrp>policy>priority-event>route-unknown

**Description** This command adds one or more route sources to match the route unknown IP route prefix for a route unknown priority control event.

If the route source does not match one of the defined protocols, the match is considered unsuccessful and the **route-unknown** event transitions to the set state.

The **protocol** command is optional. If the **protocol** command is not executed, the comparison between the RTM prefix return and the **route-unknown** IP route prefix will not include the source of the prefix. The **protocol** command cannot be executed without at least one associated route source parameter. All parameters are reset each time the **protocol** command is executed and only the explicitly defined protocols are allowed to match.

The **no** form of the command removes protocol route source as a match criteria for returned RTM route prefixes.

To remove specific existing route source match criteria, execute the **protocol** command and include only the specific route source criteria. Any unspecified route source criteria is removed.

**Default** no protocol — No route source for the route unknown priority event is defined.

- Parameters**
- bgp** — This parameter defines BGP as an eligible route source for a returned route prefix from the RTM when looking up the **route-unknown** route prefix. The **bgp** parameter is not exclusive from the other available **protocol** parameters. If **protocol** is executed without the **bgp** parameter, a returned route prefix with a source of BGP will not be considered a match and will cause the event to enter the set state. This parameter only applies to the 7750 SR and 7950 XRS.
  - bgp-vpn** — This parameter defines **bgp-vpn** as an eligible route source for a returned route prefix from the RTM when looking up the **route-unknown** route prefix. The **bgp-vpn** parameter is not exclusive from the other available **protocol** parameters. If **protocol** is executed without the **bgp-vpn** parameter, a returned route prefix with a source of **bgp-vpn** will not be considered a match and will cause the event to enter the set state. This parameter only applies to the 7750 SR and 7950 XRS.
  - ospf** — This parameter defines OSPF as an eligible route source for a returned route prefix from the RTM when looking up the **route-unknown** route prefix. The **ospf** parameter is not exclusive from the other available **protocol** parameters. If **protocol** is executed without the **ospf** parameter, a returned route prefix with a source of OSPF will not be considered a match and will cause the event to enter the set state.
  - is-is** — This parameter defines IS-IS as an eligible route source for a returned route prefix from the RTM when looking up the **route-unknown** route prefix. The **is-is** parameter is not exclusive from the other available **protocol** parameters. If **protocol** is executed without the **is-is** parameter, a returned route prefix with a source of IS-IS will not be considered a match and will cause the event to enter the set state.
  - rip** — This parameter defines RIP as an eligible route source for a returned route prefix from the RTM when looking up the **route-unknown** route prefix. The **rip** parameter is not exclusive from the other available **protocol** parameters. If **protocol** is executed without the **rip** parameter, a returned route prefix with a source of RIP will not be considered a match and will cause the event to enter the set state.
  - static** — This parameter defines a static route as an eligible route source for a returned route prefix from the RTM when looking up the **route-unknown** route prefix. The **static** parameter is not exclusive from the other available **protocol** parameters. If **protocol** is executed without the **static** parameter, a returned route prefix with a source of static route will not be considered a match and will cause the event to enter the set state.

## route-unknown

- Syntax** `[no] route-unknown [ip-prefix/mask | ipv6-address / prefix-length]`
- Context** `config>vrrp>policy>priority-event`
- Description** This command creates a context to configure a route unknown priority control event that monitors the existence of a specific active IP route prefix within the routing table.



The **route-unknown** command configures a priority control event that defines a link between the VRRP priority control policy and the Route Table Manager (RTM). The RTM registers the specified route prefix as monitored by the policy. If any change (add, delete, new next hop) occurs relative to the prefix, the policy is notified and takes correct action according to the priority event definition. If the route prefix exists and is active in the routing table according to the conditions defined, the event is in the cleared state. If the route prefix is removed, becomes inactive or fails to meet the event criteria, the event is in the set state.

The command creates a **route-unknown** node identified by *prefix/mask-length* and containing event control commands.

Multiple unique (different *prefix/mask-length*) **route-unknown** event nodes can be configured within the **priority-event** node up to the maximum limit of 32 events.

The **route-unknown** command can reference any valid IP address mask-length pair. The IP address and associated mask length define a unique IP router prefix. The dynamic monitoring of the route prefix results in one of the event operational states listed in [Table 41](#).

**Table 41**      **Route-unknown Operational States**

| route-unknown Operational State | Description                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------|
| Set – non-existent              | The route does not exist in the route table.                                                                                |
| Set – inactive                  | The route exists in the route table but is not being used.                                                                  |
| Set – wrong next hop            | The route exists in the route table but does not meet the <b>next-hop</b> requirements.                                     |
| Set – wrong protocol            | The route exists in the route table but does not meet the <b>protocol</b> requirements.                                     |
| Set – less specific found       | The route exists in the route table but does is not an exact match and does not meet any <b>less-specific</b> requirements. |
| Set – default best match        | The route exists in the route table as the default route but the default route is not allowed for route matching.           |
| Cleared – less specific found   | A less specific route exists in the route table and meets all criteria including the <b>less-specific</b> requirements.     |
| Cleared – found                 | The route exists in the route table manager and meets all criteria.                                                         |

An existing route prefix in the RTM must be active (used by the IP forwarding engine) to clear the event operational state. It may be less specific (the defined prefix may be contained in a larger prefix according to Classless Inter-Domain Routing (CIDR) techniques) if the event has the **less-specific** statement defined. The less specific route that incorporates the router prefix may be the default route (0.0.0.0) if the **less-specific allow-default** statement is defined. The matching prefix may be required to have a specific next hop IP address if defined

by the event **next-hop** command. Finally, the source of the RTM prefix may be required to be one of the dynamic routing protocols or be statically defined if defined by the event **protocol** command. If an RTM prefix is not found that matches all the above criteria (if defined in the event control commands), the event is considered to be set. If a matching prefix is found in the RTM, the event is considered to be cleared.

When an event transitions from clear to set, the set is processed immediately and must be reflected in the associated virtual router instances in-use priority value. As the event transitions from clear to set, a hold-set timer is loaded with the value configured by the events **hold-set** command. This timer prevents the event from clearing until it expires, damping the effect of event flapping. If the event clears and becomes set again before the hold-set timer expires, the timer is reset to the **hold-set** value, extending the time before another clear can take effect.

The **no** form of the command is used to remove the specific *prefix/mask-length* monitoring event. The event can be removed at anytime. When the event is removed, the in-use priority of all associated virtual router instances must be reevaluated. The events **hold-set** timer has no effect on the removal procedure.

**Default** no route-unknown — No route unknown priority control events are defined for the priority control event policy.

**Parameters** *prefix* — The IP prefix address to be monitored by the route unknown priority control event in dotted decimal notation.

**Values** 0.0.0.0 to 255.255.255.255

*mask-length* — The subnet mask length expressed as a decimal integer associated with the IP *prefix* defining the route prefix to be monitored by the route unknown priority control event.

**Values** 0 to 32

*ip-address* — The IP address of the host for which the specific event will monitor connectivity. The *ip-addr* can only be monitored by a single event in this policy. The IP address can be monitored by multiple VRRP priority control policies. The IP address can be used in one or multiple **ping** requests. Each VRRP priority control **host-unreachable** and **ping** destined to the same *ip-addr* is uniquely identified on a per message basis. Each session originates a unique identifier value for the ICMP echo request messages it generates. This allows received ICMP echo reply messages to be directed to the appropriate sending application.

**Values** The following values apply to the 7450 ESS:

|                   |           |                            |
|-------------------|-----------|----------------------------|
| <i>ip-prefix/</i> | ip-prefix | a.b.c.d (host bits must be |
| <i>mask:</i>      |           | 0)                         |
|                   | mask      | 0 to 32                    |

**Values** The following values apply to the 7750 SR and 7950 XRS:

|                        |           |                               |
|------------------------|-----------|-------------------------------|
| <i>ip-prefix/mask:</i> | ip-prefix | a.b.c.d (host bits must be 0) |
|                        | mask      | 0 to 32                       |

---

*ipv6-address/prefix:* ipv6-address x:x:x:x:x:x:x (eight 16-bit pieces)  
x:x:x:x:x:d.d.d.d  
x: [0..FFFF]H  
prefix-length 1 to 128



## 3.10 Show, Monitor, Clear, and Debug Command Reference

- [Command Hierarchies](#)
- [Command Descriptions](#)

The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

### 3.10.1 Command Hierarchies

- [Show Commands](#)
- [Monitor Commands](#)
- [Clear Commands](#)
- [Debug Commands](#)

#### 3.10.1.1 Show Commands

```
show
 — vrrp
 — policy [policy-id [event event-type specific-qualifier]]
 — router
 — vrrp
 — instance
 — instance [interface interface-name [vrid virtual-router-id]]
 — instance interface interface-name vrid virtual-router-id ipv6
 — statistics
```

#### 3.10.1.2 Monitor Commands

```
monitor
 — router
 — vrrp
 — instance interface interface-name vr-id virtual-router-id [ipv6] [interval
 seconds] [repeat repeat] [absolute | rate]
```

### 3.10.1.3 Clear Commands

```
clear
 — vrrp
 — statistics
 — router
 — vrrp
 — interface ip-int-name [vrid virtual-router-id]
 — interface ip-int-name vrid virtual-router-id ipv6
 — statistics interface interface-name [vrid virtual-router-id]
 — statistics
 — statistics interface interface-name vrid virtual-router-id ipv6
```

### 3.10.1.4 Debug Commands

```
debug
 — router
 — vrrp
 — events
 — events interface ip-int-name [vrid virtual-router-id]
 — events interface ip-int-name vrid virtual-router-id ipv6
 — no events
 — no events interface ip-int-name [vrid virtual-router-id]
 — no events interface ip-int-name vrid virtual-router-id ipv6
 — packets
 — packets interface ip-int-name [vrid virtual-router-id]
 — packets interface ip-int-name vrid virtual-router-id ipv6
 — no packets
 — no packets interface ip-int-name [vrid virtual-router-id]
 — no packets interface ip-int-name vrid virtual-router-id ipv6
```

## 3.10.2 Command Descriptions

### 3.10.2.1 Show Commands

The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

instance

**Syntax**    **instance**

**instance** [**interface** *interface-name* [**vrid** *virtual-router-id*]]  
**instance interface** *interface-name* **vrid** *virtual-router-id* **ipv6**

**Context** show>vrrp

**Description** This command displays information for VRRP instances.

If no command line options are specified, summary information for all VRRP instances displays.

**Parameters** **interface** *ip-int-name* — Displays detailed information for the VRRP instances on the specified IP interface including status and statistics.

**Default** Summary information for all VRRP instances.

**vrid** *virtual-router-id* — Displays detailed information for the specified VRRP instance on the IP interface.

**Default** All VRIDs for the IP interface.

**Values** 1 to 255

**ipv6** — Specifies the IPv6 instance.

**Output** The following output is an example of VRRP instance information for the 7450 ESS, and [Table 42](#) describes the fields.

### Sample Output

```
*A:ALA-A# show router vrrp instance
=====
VRRP Instances
=====
Interface Name VR Id Own Adm State Base Pri Msg Int
 IP Opr Pol Id InUse Pri Inh Int

n2 1 No Up Master 100 1
 IPv4 Up n/a 100 No
 Backup Addr: 5.1.1.10

Instances : 2
=====
*A:ALA-A#

*A:ALA-A# show router vrrp instance interface n2 vrid 1
=====
VRRP Instance 1 for interface "n2"
=====
Owner : No VRRP State : Master
Primary IP of Master: 5.1.1.2 (Self)
Primary IP : 5.1.1.2 Standby-Forwarding: Disabled
VRRP Backup Addr : 5.1.1.10
Admin State : Up Oper State : Up
Up Time : 09/23/2004 06:53:45 Virt MAC Addr : 00:00:5e:00:01:01
Auth Type : None
Config Mesg Intvl : 1 In-Use Mesg Intvl : 1
```

```

Master Inherit Intvl: No
Base Priority : 100
Policy ID : n/a
Ping Reply : No
SSH Reply : No
Init Delay : 0
Creation State : Active
OperGroup : vrrp1_1
In-Use Priority : 100
Preempt Mode : Yes
Telnet Reply : No
Traceroute Reply : No
Init Timer Expires: 0.000 sec

```

#### ----- Master Information

```

Primary IP of Master: 5.1.1.2 (Self)
Addr List Mismatch : No
Master Since : 09/23/2004 06:53:49
Master Priority : 100

```

#### ----- Masters Seen (Last 32)

| Primary IP of Master | Last Seen           | Addr List Mismatch | Msg Count |
|----------------------|---------------------|--------------------|-----------|
| 5.1.1.2              | 09/23/2004 06:53:49 | No                 | 0         |

#### ----- Statistics

|                     |       |                    |     |
|---------------------|-------|--------------------|-----|
| Become Master       | : 1   | Master Changes     | : 1 |
| Adv Sent            | : 103 | Adv Received       | : 0 |
| Pri Zero Pkts Sent  | : 0   | Pri Zero Pkts Rcvd | : 0 |
| Preempt Events      | : 0   | Preempted Events   | : 0 |
| Mesg Intvl Discards | : 0   | Mesg Intvl Errors  | : 0 |
| Addr List Discards  | : 0   | Addr List Errors   | : 0 |
| Auth Type Mismatch  | : 0   | Auth Failures      | : 0 |
| Invalid Auth Type   | : 0   | Invalid Pkt Type   | : 0 |
| IP TTL Errors       | : 0   | Pkt Length Errors  | : 0 |
| Total Discards      | : 0   |                    |     |

The following output is an example of VRRP instance information for the 7750 SR and 7950 XRS, and [Table 42](#) describes the fields

### Output Sample

```

*A:ALA-A# show router vrrp instance interface n2 vrid 1 ipv6
=====
VRRP Instance 1 for interface "n2"
=====
No Matching Entries
=====
*A:ALA-A#

*A:ALA-A# show router vrrp instance interface n2 vrid 10 ipv6
=====
VRRP Instance 10 for interface "n2"
=====
Owner : No
Primary IP of Master: FE80::1 (Self)
Primary IP : FE80::1
VRRP State : Master
Standby-Forwarding: Disabled

```



```

VRRP Backup Addr : 5::10
 : FE80::10
Admin State : Up Oper State : Up
Up Time : 09/23/2004 06:55:12 Virt MAC Addr : 00:00:5e:00:02:0a
Config Mesg Intvl : 1.0 In-Use Mesg Intvl : 1.0
Master Inherit Intvl : Yes
Base Priority : 100 In-Use Priority : 100
Policy ID : n/a Preempt Mode : Yes
Ping Reply : No Telnet Reply : No
 : Traceroute Reply : No
Init Delay : 0 Init Timer Expires : 0.000 sec
Creation State : Active

Master Information

Primary IP of Master: FE80::1 (Self)
Addr List Mismatch : No Master Priority : 100
Master Since : 09/23/2004 06:55:16
=====
Masters Seen (Last 32)
=====
Primary IP of Master Last Seen Addr List Mismatch Msg Count

FE80::1 09/23/2004 06:55:16 No 0

Statistics

Master Transitions : 1 Discontinuity Time: 09/09/2004 01:57*
Adv Sent : 23 Adv Received : 0
Pri Zero Pkts Sent : 0 Pri Zero Pkts Rcvd : 0
Preempt Events : 0 Preempted Events : 0
Mesg Intvl Discards : 0 Mesg Intvl Errors : 0
Total Discards : 0 Addr List Errors : 0
Auth Failures : 0 Invalid Pkt Type : 0
IP TTL Errors : 0 Pkt Length Errors : 0
=====
* indicates that the corresponding row element may have been truncated.

```

**Table 42** Show VRRP Instance Output Fields

| Label          | Description                                                                                  |
|----------------|----------------------------------------------------------------------------------------------|
| Interface name | The name of the IP interface.                                                                |
| VR ID          | The virtual router ID for the IP interface                                                   |
| Own<br>Owner   | Yes<br>Specifies that the virtual router instance as owning the virtual router IP addresses. |
|                | No<br>Indicates that the virtual router instance is operating as a non-owner.                |

**Table 42 Show VRRP Instance Output Fields (Continued)**

| Label          | Description                                                                                                                                                                                                                                                                                                                           |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Adm            | Up<br>Indicates that the administrative state of the VRRP instance is up.                                                                                                                                                                                                                                                             |
|                | Down<br>Indicates that the administrative state of the VRRP instance is down.                                                                                                                                                                                                                                                         |
| Opr            | Up<br>Indicates that the operational state of the VRRP instance is up.                                                                                                                                                                                                                                                                |
|                | Down<br>Indicates that the operational state of the VRRP instance is down.                                                                                                                                                                                                                                                            |
| State          | When owner, <b>backup</b> defines the IP addresses that are advertised within VRRP advertisement messages.<br>When non-owner, <b>backup</b> actually creates an IP interface IP address used for routing IP packets and communicating with the system when the access commands are defined (ping-reply, telnet-reply, and ssh-reply). |
| Pol Id         | The value that uniquely identifies a Priority Control Policy.                                                                                                                                                                                                                                                                         |
| Base Priority  | The <i>base-priority</i> value used to derive the in-use priority of the virtual router instance as modified by any optional VRRP priority control policy.                                                                                                                                                                            |
| InUse Priority | The current in-use priority associated with the VRRP virtual router instance.                                                                                                                                                                                                                                                         |
| Msg Int        | The administrative advertisement message timer used by the master virtual router instance to send VRRP advertisement messages and to derive the master down timer as backup.                                                                                                                                                          |

**Table 42 Show VRRP Instance Output Fields (Continued)**

| Label        | Description                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Inh Int      | Yes<br>When the VRRP instance is a non-owner and is operating as a backup and the <b>master-int-inherit</b> command is enabled, the master down timer is indirectly derived from the value in the advertisement interval field of the VRRP message received from the current master.                                                                                                                                                     |
|              | No<br>When the VRRP instance is operating as a backup and the <b>master-int-inherit</b> command is <i>not</i> enabled, the configured advertisement interval is matched against the value in the advertisement interval field of the VRRP message received from the current master. If the two values do not match then the VRRP advertisement is discarded.<br>If the VRRP instance is operating as a master, this value has no effect. |
| Backup Addr  | The backup virtual router IP address.                                                                                                                                                                                                                                                                                                                                                                                                    |
| BFD          | Indicates BFD is enabled.                                                                                                                                                                                                                                                                                                                                                                                                                |
| VRRP State   | Specifies whether the VRRP instance is operating in a master or backup state.                                                                                                                                                                                                                                                                                                                                                            |
| Policy ID    | The VRRP priority control policy associated with the VRRP virtual router instance.<br>A value of 0 indicates that no control policy policy is associated with the virtual router instance.                                                                                                                                                                                                                                               |
| Preempt Mode | Yes<br>The preempt mode is enabled on the virtual router instance where it will preempt a VRRP master with a lower priority.                                                                                                                                                                                                                                                                                                             |
|              | No<br>The preempt mode is disabled and prevents the non-owner virtual router instance from preempting another, less desirable virtual router.                                                                                                                                                                                                                                                                                            |

**Table 42 Show VRRP Instance Output Fields (Continued)**

| Label                | Description                                                                                                                                                                                                                                                                                                                                                |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Ping Reply           | Yes<br>A non-owner master is enabled to reply to ICMP Echo requests directed to the virtual router instance IP addresses.<br>Ping Reply is valid only if the VRRP virtual router instance associated with this entry is a non-owner.<br>A non-owner backup virtual router never responds to such ICMP echo requests irrespective if Ping Reply is enabled. |
|                      | No<br>ICMP echo requests to the virtual router instance IP addresses are discarded.                                                                                                                                                                                                                                                                        |
| Telnet Reply         | Yes<br>Non-owner masters can to reply to TCP port 23 Telnet requests directed at the virtual router instances IP addresses.                                                                                                                                                                                                                                |
|                      | No<br>Telnet requests to the virtual router instance IP addresses are discarded.                                                                                                                                                                                                                                                                           |
| SSH Reply            | Yes<br>Non-owner masters can to reply to SSH requests directed at the virtual router instances IP addresses.                                                                                                                                                                                                                                               |
|                      | No<br>All SSH request messages destined to the non-owner virtual router instance IP addresses are discarded.                                                                                                                                                                                                                                               |
| Primary IP of Master | The IP address of the VRRP master.                                                                                                                                                                                                                                                                                                                         |
| Primary IP           | The IP address of the VRRP owner.                                                                                                                                                                                                                                                                                                                          |
| Up Time              | The date and time when the operational state of the event last changed.                                                                                                                                                                                                                                                                                    |
| Virt MAC Addr        | The virtual MAC address used in ARP responses when the VRRP virtual router instance is operating as a master.                                                                                                                                                                                                                                              |
| Auth Type            | Specifies the VRRP authentication Type 0 (no authentication), Type 1 (simple password), or Type 2 (MD5) for the virtual router.                                                                                                                                                                                                                            |
| Addr List Mismatch   | Specifies whether a trap was generated when the IP address list received in the advertisement messages received from the current master did not match the configured IP address list.<br>This is an edge triggered notification. A second trap will not be generated for a packet from the same master until this event has been cleared.                  |

**Table 42 Show VRRP Instance Output Fields (Continued)**

| Label           | Description                                                                                                                                                                                                                                                         |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Master Priority | The priority of the virtual router instance which is the current master.                                                                                                                                                                                            |
| Master Since    | The date and time when operational state of the virtual router changed to master.<br>For a backup virtual router, this value specifies the date and time when it received the first VRRP advertisement message from the virtual router which is the current master. |
| OperGroup       | Displays the operational group name associated with the VRRP interface, if configured.                                                                                                                                                                              |

## policy

**Syntax** **policy** [*vrrp-policy-id* [**event** *event-type specific-qualifier*]]

**Context** show>vrrp

**Description** This command displays VRRP priority control policy information.

If no command line options are specified, a summary of the VRRP priority control event policies displays.

**Parameters** *vrrp-policy-id* — Displays information on the specified priority control policy ID.

**Default** All VRRP policies IDs

**Values** 1 to 9999

**event event-type** — Displays information on the specified VRRP priority control event within the policy ID.

**Default** All event types and qualifiers

**Values** **port-down** *port-id*  
**lag-port-down** *lag-id*  
**host-unreachable** *host-ip-addr*  
**route-unknown** *route-prefix/mask*  
**mc-ipsec-non-forwarding**

*specific-qualifier* — Display information about the specified qualifier.

**Values** port-id, lag-id, host-ip-addr, route-prefix/mask, tunnel-group-id

**Output** **VRRP Policy Output** — The following output is an example of VRRP policy information, and [Table 43](#) describes the fields.

**Table 43**      **Show VRRP Policy Output Fields**

| Label                      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Policy Id                  | The VRRP priority control policy associated with the VRRP virtual router instance.<br><br>A value of 0 indicates that no control policy is associated with the virtual router instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Current Priority & Effects |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Current Explicit           | When multiple explicitly defined events associated with the priority control policy happen simultaneously, the lowest value of all the current explicit priorities will be used as the in-use priority for the virtual router.                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Current Delta Sum          | The sum of the priorities of all the delta events when multiple delta events associated with the priority control policy happen simultaneously. This sum is subtracted from the base priority of the virtual router to give the in-use priority.                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Delta Limit                | The delta-in-use-limit for a VRRP policy. Once the total sum of all delta events has been calculated and subtracted from the base-priority of the virtual router, the result is compared to the delta-in-use-limit value. If the result is less than this value, the delta-in-use-limit value is used as the virtual router in-use priority value. If an explicit priority control event overrides the delta priority control events, the delta-in-use-limit has no effect.<br><br>If the delta-in-use-limit is 0, the sum of the delta priority control events to reduce the virtual router's in-use-priority to 0 can prevent it from becoming or staying master. |
| Current Priority           | The configured delta-in-use-limit priority for a VRRP priority control policy or the configured delta or explicit priority for a priority control event.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Applied                    | The number of virtual router instances to which the policy has been applied. The policy cannot be deleted unless this value is 0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Description                | A text string which describes the VRRP policy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

**Table 43 Show VRRP Policy Output Fields (Continued)**

| Label              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Event Type & ID    | <p>A delta priority event is a conditional event defined in a priority control policy that subtracts a given amount from the base priority to give the current in-use priority for the VRRP virtual router instances to which the policy is applied.</p> <p>An explicit priority event is a conditional event defined in a priority control policy that explicitly defines the in-use priority for the VRRP virtual router instances to which the policy is applied.</p> <p>Explicit events override all delta Events. When multiple explicit events occur simultaneously, the event with the lowest priority value defines the in-use priority.</p>                                                                                                                                                                                                                                                                                                                                               |
| Event Oper State   | The operational state of the event.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Hold Set Remaining | The amount of time that must pass before the set state for a VRRP priority control event can transition to the cleared state to dampen flapping events.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Priority & Effect  | <p><b>Delta</b></p> <p>The <i>priority-level</i> value is subtracted from the associated virtual router instance's base priority when the event is set and no explicit events are set. The sum of the priority event <i>priority-level</i> values on all set delta priority events are subtracted from the virtual router base priority to derive the virtual router instance in-use priority value.</p> <p>If the <b>delta</b> priority event is cleared, the <i>priority-level</i> is no longer used in the in-use priority calculation.</p> <p><b>Explicit</b></p> <p>The <i>priority-level</i> value is used to override the base priority of the virtual router instance if the priority event is set and no other <b>explicit</b> priority event is set with a lower <i>priority-level</i>.</p> <p>The set <b>explicit</b> priority value with the lowest <i>priority-level</i> determines the actual in-use protocol value for all virtual router instances associated with the policy.</p> |
| In Use             | Specifies whether the event is currently affecting the in-use priority of some virtual router.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**Sample Output**

```

A:ALA-A# show vrrp policy
=====
VRRP Policies
=====
Policy Current Current Current Delta Applied
Id Priority & Effect Explicit Delta Sum Limit

```

```

1 None None None 1 Yes
2 None None None 1 No
=====
A:ALA-A#

A:ALA-A# show vrrp policy 1
=====
VRRP Policy 1
=====
Description : 10.10.200.253 reachability
Current Priority: None Applied : No
Current Explicit: None Current Delta Sum : None
Delta Limit : 1

Applied To VR Opr Base In-use Master Is
Interface Name Id Pri Pri Pri Pri Master

None

Priority Control Events

Event Type & ID Event Oper State Hold Set Priority In
 Remaining &Effect Use

Host Unreach 10.10.200.252 n/a Expired 20 Del No
Host Unreach 10.10.200.253 n/a Expired 10 Del No
Route Unknown 10.10.100.0/24 n/a Expired 1 Exp No
=====

```

**VRRP Policy Event Output** — The following output is an example of VRRP policy event information, and [Table 44](#) describes the fields.

### Sample Output

```

A:ALA-A#show vrrp policy 1 event port-down
=====
VRRP Policy 1, Event Port Down 1/1/1
=====
Description :
Current Priority: None Applied : Yes
Current Explicit: None Current Delta Sum : None
Delta Limit : 1

Applied To VR Opr Base In-use Master Is
Interface Name Id Pri Pri Pri Pri Master

ies301backup 1 Down 100 100 0 No

Priority Control Event Port Down 1/1/1

Priority : 30 Priority Effect : Delta

```



```
Hold Set Config : 0 sec Hold Set Remaining: Expired
Value In Use : No Current State : Cleared
trans to Set : 6 Previous State : Set-down
Last Transition : 04/13/2007 04:54:35
```

=====

A:ALA-A#

A:ALA-A# show vrrp policy 1 event host-unreachable

=====

VRRP Policy 1, Event Host Unreachable 10.10.200.252

=====

```
Description : 10.10.200.253 reachability
Current Priority: None Applied : No
Current Explicit: None Current Delta Sum : None
Delta Limit : 1
```

```

Applied To VR Opr Base In-use Master Is
Interface Name Id Pri Pri Pri Master

None
```

-----

Priority Control Event Host Unreachable 10.10.200.252

-----

```
Priority : 20 Priority Effect : Delta
Interval : 1 sec Timeout : 1 sec
Drop Count : 3
Hold Set Config : 0 sec Hold Set Remaining: Expired
Value In Use : No Current State : n/a
trans to Set : 0 Previous State : n/a
Last Transition : 04/13/2007 23:10:24
```

=====

A:ALA-A#

A:ALA-A# show vrrp policy 1 event route-unknown

=====

VRRP Policy 1, Event Route Unknown 10.10.100.0/24

=====

```
Description : 10.10.200.253 reachability
Current Priority: None Applied : No
Current Explicit: None Current Delta Sum : None
Delta Limit : 1
```

```

Applied To VR Opr Base In-use Master Is
Interface Name Id Pri Pri Pri Master

None
```

-----

Priority Control Event Route Unknown 10.10.100.0/24

-----

```
Priority : 1 Priority Effect : Explicit
Less Specific : No Default Allowed : No
Next Hop(s) : None
Protocol(s) : None
Hold Set Config : 0 sec Hold Set Remaining: Expired
```

```

Value In Use : No Current State : n/a
trans to Set : 0 Previous State : n/a
Last Transition : 04/13/2007 23:10:24
=====

```

**Table 44** Show VRRP Policy Event Output Fields

| Label                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Description               | A text string which describes the VRRP policy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Policy Id                 | The VRRP priority control policy associated with the VRRP virtual router instance.<br>A value of 0 indicates that no control policy is associated with the virtual router instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Current Priority          | The base router priority for the virtual router instance used in the master election process.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Current Explicit          | When multiple explicitly defined events associated with the priority control policy happen simultaneously, the lowest value of all the current explicit priorities will be used as the in-use priority for the virtual router.                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Applied                   | The number of virtual router instances to which the policy has been applied. The policy cannot be deleted unless this value is 0.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Current Delta Sum         | The sum of the priorities of all the delta events when multiple delta events associated with the priority control policy happen simultaneously. This sum is subtracted from the base priority of the virtual router to give the in-use priority.                                                                                                                                                                                                                                                                                                                                                                                                                |
| Delta Limit               | The delta-in-use-limit for a VRRP policy. Once the total sum of all delta events has been calculated and subtracted from the base-priority of the virtual router, the result is compared to the delta-in-use-limit value. If the result is less than this value, the delta-in-use-limit value is used as the virtual router in-use priority value. If an explicit priority control event overrides the delta priority control events, the delta-in-use-limit has no effect.<br>If the delta-in-use-limit is 0, the sum of the delta priority control events to reduce the virtual router's in-use-priority to 0 can prevent it from becoming or staying master. |
| Applied to Interface Name | The interface name where the VRRP policy is applied.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| VR ID                     | The virtual router ID for the IP interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**Table 44 Show VRRP Policy Event Output Fields (Continued)**

| Label              | Description                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Opr                | Up<br>Indicates that the operational state of the VRRP instance is up.                                                                                                                                                                                                                                                                                         |
|                    | Down<br>Indicates that the operational state of the VRRP instance is down.                                                                                                                                                                                                                                                                                     |
| Base Pri           | The base priority used by the virtual router instance.                                                                                                                                                                                                                                                                                                         |
| InUse Priority     | The current in-use priority associated with the VRRP virtual router instance.                                                                                                                                                                                                                                                                                  |
| Master Priority    | The priority of the virtual router instance which is the current master.                                                                                                                                                                                                                                                                                       |
| Priority           | The base priority used by the virtual router instance.                                                                                                                                                                                                                                                                                                         |
| Priority Effect    | Delta<br>A delta priority event is a conditional event defined in a priority control policy that subtracts a given amount from the base priority to give the current in-use priority for the VRRP virtual router instances to which the policy is applied.                                                                                                     |
|                    | Explicit<br>A conditional event defined in a priority control policy that explicitly defines the in-use priority for the VRRP virtual router instances to which the policy is applied.<br>Explicit events override all delta events. When multiple explicit events occur simultaneously, the event with the lowest priority value defines the in-use priority. |
| Current Priority   | The configured delta-in-use-limit priority for a VRRP priority control policy or the configured delta or explicit priority for a priority control event.                                                                                                                                                                                                       |
| Event Oper State   | The operational state of the event.                                                                                                                                                                                                                                                                                                                            |
| Hold Set Remaining | The amount of time that must pass before the set state for a VRRP priority control event can transition to the cleared state to dampen flapping events.                                                                                                                                                                                                        |
| Priority           | The base priority used by the virtual router instance.                                                                                                                                                                                                                                                                                                         |

**Table 44 Show VRRP Policy Event Output Fields (Continued)**

| Label           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Priority Effect | <p><b>Delta</b></p> <p>The <i>priority-level</i> value is subtracted from the associated virtual router instance's base priority when the event is set and no explicit events are set. The sum of the priority event <i>priority-level</i> values on all set delta priority events are subtracted from the virtual router base priority to derive the virtual router instance in-use priority value.</p> <p>If the <b>delta</b> priority event is cleared, the <i>priority-level</i> is no longer used in the in-use priority calculation.</p> |
|                 | <p><b>Explicit</b></p> <p>The <i>priority-level</i> value is used to override the base priority of the virtual router instance if the priority event is set and no other <b>explicit</b> priority event is set with a lower <i>priority-level</i>.</p> <p>The set <b>explicit</b> priority value with the lowest <i>priority-level</i> determines the actual in-use protocol value for all virtual router instances associated with the policy.</p>                                                                                            |
| Hold Set Config | The configured number of seconds that the hold-set timer waits after an event enters a set state or enters a higher threshold set state, depending on the event type.                                                                                                                                                                                                                                                                                                                                                                          |
| Value In Use    | <p><b>Yes</b></p> <p>The event is currently affecting the in-use priority of some virtual router.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|                 | <p><b>No</b></p> <p>The event is not affecting the in-use priority of some virtual router.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| # trans to Set  | The number of times the event has transitioned to one of the 'set' states.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Last Transition | The time and date when the operational state of the event last changed.                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

## statistics

**Syntax**     **statistics**

**Context**    show>router>vrrp

**Description**    This command displays statistics for VRRP instance.

**Output**        The following output is an example of VRRP statistics information, and table describes the fields.

**Sample Output**

```
A:ALA-48# show router vrrp statistics
=====
VRRP Global Statistics
=====
VR Id Errors : 0 Version Errors : 0
Checksum Errors : 0
=====
```

**Table 45**      **Show VRRP Statistics Output Fields**

| Label           | Description                                      |
|-----------------|--------------------------------------------------|
| VR Id Errors    | Displays the number of virtual router ID errors. |
| Version Errors  | Displays the number of version errors.           |
| Checksum Errors | Displays the number of checksum errors.          |

**3.10.2.2    Monitor Commands**

instance

- Syntax**      **instance interface** *interface-name* **vr-id** *virtual-router-id* [**ipv6**] [**interval seconds**] [**repeat repeat**] [**absolute | rate**]
- Context**      monitor>router>vrrp
- Description**      Monitor statistics for a VRRP instance.
- Parameters**      *interface-name* — The name of the existing IP interface on which VRRP is configured.  
**vr-id** *virtual-router-id* — The virtual router ID for the existing IP interface, expressed as a decimal integer.  
**interval seconds** — Configures the interval for each display in seconds.  
                  **Values**      3 to 60  
                  **Default**    5 seconds  
**repeat repeat** — Configures how many times the command is repeated.  
                  **Values**      1 to 999  
                  **Default**    10  
**absolute** — When the **absolute** keyword is specified, the raw statistics are displayed, without processing. No calculations are performed on the delta or rate statistics.

**rate** — When the **rate** keyword is specified, the rate-per-second for each statistic is displayed instead of the delta.

**ipv6** — Specifies to monitor IPv6 instances.

**Output** The following output is an example of VRRP instance information.

### Sample Output

```
*A:ALA-A# monitor router vrrp instance interface n2 vr-id 1
=====
Monitor statistics for VRRP Instance 1 on interface "n2"
=====

At time t = 0 sec (Base Statistics)

Become Master : 1 Master Changes : 1
Adv Sent : 1439 Adv Received : 0
Pri Zero Pkts Sent : 0 Pri Zero Pkts Rcvd : 0
Preempt Events : 0 Preempted Events : 0
Mesg Intvl Discards : 0 Mesg Intvl Errors : 0
Addr List Discards : 0 Addr List Errors : 0
Auth Type Mismatch : 0 Auth Failures : 0
Invalid Auth Type : 0 Invalid Pkt Type : 0
IP TTL Errors : 0 Pkt Length Errors : 0
Total Discards : 0
=====
```

The following output is an example of VRRP instance information for the 7750 SR and 7950 XRS.

### Sample Output

```
*A:ALA-A# monitor router vrrp instance interface n2 vr-id 10 ipv6
=====
Monitor statistics for VRRP Instance 10 on interface "n2"
=====

At time t = 0 sec (Base Statistics)

Master Transitions : 1 Discontinuity Time: 09/09/2004 01:57*
Adv Sent : 1365 Adv Received : 0
Pri Zero Pkts Sent : 0 Pri Zero Pkts Rcvd : 0
Preempt Events : 0 Preempted Events : 0
Mesg Intvl Discards : 0 Mesg Intvl Errors : 0
Total Discards : 0 Addr List Errors : 0
Auth Failures : 0 Invalid Pkt Type : 0
IP TTL Errors : 0 Pkt Length Errors : 0
=====
```

### 3.10.2.3 Clear Commands

#### interface

|                    |                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>interface</b> <i>ip-int-name</i> [ <b>vrid</b> <i>virtual-router-id</i> ]<br><b>interface</b> <i>ip-int-name</i> <b>vrid</b> <i>virtual-router-id</i> <b>ipv6</b>                                                                                                                                                                    |
| <b>Context</b>     | clear>router>vrrp                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b> | This command resets VRRP protocol instances on an IP interface.                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <i>ip-int-name</i> — The IP interface to reset the VRRP protocol instances.<br><b>vrid</b> <i>vrid</i> — Resets the VRRP protocol instance for the specified VRID on the IP interface.<br><b>Default</b> All VRIDs on the IP interface.<br><b>Values</b> 1 to 255<br><b>ipv6</b> — Clears IPv6 information for the specified interface. |

#### statistics

|                    |                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>statistics</b> [ <b>policy</b> <i>policy-id</i> ]                                                    |
| <b>Context</b>     | clear>router>vrrp                                                                                       |
| <b>Description</b> | This command enables the context to clear and reset VRRP entities.                                      |
| <b>Parameters</b>  | <b>policy</b> <i>policy-id</i> — Clears statistics for the specified policy.<br><b>Values</b> 1 to 9999 |

#### statistics

|                    |                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>statistics interface</b> <i>interface-name</i> [ <b>vrid</b> <i>virtual-router-id</i> ]<br><b>statistics</b><br><b>statistics interface</b> <i>interface-name</i> <b>vrid</b> <i>virtual-router-id</i> <b>ipv6</b> |
| <b>Context</b>     | clear>router>vrrp                                                                                                                                                                                                     |
| <b>Description</b> | This command clears statistics for VRRP instances on an IP interface or VRRP priority control policies.                                                                                                               |
| <b>Parameters</b>  | <b>interface</b> <i>ip-int-name</i> — Clears the VRRP statistics for all VRRP instances on the specified IP interface.                                                                                                |

**vrid** *virtual-router-id* — Clears the VRRP statistics for the specified VRRP instance on the IP interface.

**Default** All VRRP instances on the IP interface.

**Values** 1 to 255

**policy** [*vrrp-policy-id*] — Clears VRRP statistics for all or the specified VRRP priority control policy.

**Default** All VRRP policies.

**Values** 1 to 9999

**ipv6** — Clears IPv6 statistics for the specified interface.

### 3.10.2.4 Debug Commands

#### events

**Syntax** **events**  
**events interface** *ip-int-name* [**vrid** *virtual-router-id*]  
**events interface** *ip-int-name* **vrid** *virtual-router-id* **ipv6**  
**no events**  
**no events interface** *ip-int-name* **vrid** *virtual-router-id* **ipv6**  
**no events interface** *ip-int-name* [**vrid** *virtual-router-id*]

**Context** debug>router>vrrp

**Description** This command enables debugging for VRRP events.

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

**Parameters** *ip-int-name* — Displays the specified interface name.

**vrid** *virtual-router-id* — Displays the specified VRID.

**ipv6** — Debugs the specified IPv6 VRRP interface.

#### packets

**Syntax** **packets interface** *ip-int-name* [**vrid** *virtual-router-id*]  
**packets**  
**no packets interface** *ip-int-name* [**vrid** *virtual-router-id*] [**ipv6**]  
**no packets**

**Context** debug>router>vrrp



---

|                    |                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | This command enables debugging for VRRP packets.<br><br>The <b>no</b> form of the command disables debugging.                      |
| <b>Parameters</b>  | <i>ip-int-name</i> — Displays the specified interface name.<br><b>vrid</b> <i>virtual-router-id</i> — Displays the specified VRID. |



---

## 4 Filter Policies

### 4.1 ACL Filter Policy Overview

ACL filter policies, also referred to as Access Control Lists (ACLs) or just “filters”, are sets of ordered rule entries specifying packet match criteria and actions to be performed to a packet upon a match. Filter policies are created with a unique filter ID, but each filter can also have a unique filter name configured after the filter policy has been created. Either filter ID or filter name can be used throughout the system to manage filter policies and assign them to interfaces.

There are three main types of filter policies: IPv4, IPv6, and MAC filter policies. MAC filter policies support three sub-types: **configure>filter>mac-filter>type {normal | isid | vid}**. These sub-types allow different Layer 2 match criteria for a MAC filter to be configured.

Additionally, nodes that support Network Group Encryption (NGE) can use IP exception filters. IP exception filters scan all outbound traffic entering an NGE domain and allow packets that match the exception filter criteria to transit the NGE domain unencrypted.

There are different kinds of filter policies as defined by the filter policy **scope**:

- An **exclusive** filter defines policy rules explicitly for a single interface. An exclusive filter allows the highest level of customization but uses the most resources, because each exclusive filter consumes hardware resources on line cards on which the interface exists.
- A **template** filter uses an identical set of policy rules across multiple interfaces. Template filters use a single set of resources per line card, regardless of how many interfaces use a specific template filter policy on that line card. Template filter policies used on access interfaces consume resources on line cards only if at least one access interface for a specific template filter policy is configured on a specific line card.
- An **embedded** filter defines a common set of policy rules that can then be used (embedded) by other exclusive or template filters in the system. This allows optimized management of filter policies.

- A **system** filter policy defines a common set of policy rules that can then be activated within other exclusive/template filters. A system filter policy is intended mainly for system-level blacklisting rules but can be used for other applications as well. This allows optimized management of common rules (similarly to embedded filters). However, active system filter policy entries are not duplicated inside each policy that activates the system policy (as is the case when embedding is used). The active system policy is downloaded once to line cards, and activated filter policies are chained to it.

Once created, filter policies must then be associated with interfaces/services/subscribers or with other filter policies (if the created policy cannot be directly deployed on an interface/service/subscriber), so the incoming/outgoing traffic can be subjected to filter rules. Filter policies are associated with interfaces/services/subscribers separately in the ingress and egress directions. A policy deployed on ingress and egress direction can be the same or different. In general, Nokia recommends using different filter policies for the ingress and egress directions and to use different filter policies per service type, since filter policies support different match criteria and different actions for different directions/service contexts.

A filter policy is applied to a packet in the ascending rule entry order. When a packet matches all the parameters specified in a filter entry's match criteria, the system takes the action defined for that entry. If a packet does not match the entry parameters, the packet is compared to the next higher numerical filter entry rule, and so on. If the packet does not match any of the entries, the system executes the **default-action** specified in the filter policy: **drop** or **forward**.

For Layer 2, either an IPv4/IPv6 or MAC filter policy can be applied. For Layer 3 and network interfaces, an IPv4/IPv6 policy can be applied. For R-VPLS service, a Layer 2 filter policy can be applied to Layer 2 forwarded traffic and a Layer 3 filter policy can be applied to Layer 3 routed traffic. For dual-stack interfaces, if both IPv4 and IPv6 filter policies are configured, the policy applied will be based on the outer IP header of the packet. Non-IP packets do not affect an IP filter policy, so the default action in the IP filter policy will not apply to these packets. IPv6 filters do not apply to the 7450 ESS except when it is in mixed mode.

### 4.1.1 Filter Policy Basics

The following subsections define main functionality supported by filter policies.

### 4.1.1.1 Filter Policy Packet Match Criteria

This section defines packet match criteria supported on SR OS for IPv4, IPv6, and MAC filters. Supported criteria types depend on the hardware platform and filter direction, see your Nokia representative for more information.

General notes:

- If multiple unique match criteria are specified in a single filter policy entry, all criteria must be met in order for the packet to be considered a match against that filter policy entry (logical AND).
- Any match criteria not explicitly defined is ignored during match.
- An ACL filter policy entry with match criteria defined, but no action configured, is considered incomplete and inactive (an entry is not downloaded to the line card). A filter policy must have at least one entry active for the policy to be considered active.
- An ACL filter entry with no match conditions defined matches all packets.
- Because an ACL filter policy is an ordered list, entries should be configured (numbered) from the most explicit to the least explicit.

### 4.1.1.2 IPv4/IPv6 Filter Policy Entry Match Criteria

The IPv4 and IPv6 match criteria supported by SR OS are listed below. The criteria are evaluated against the outer IPv4/IPv6 header and a Layer 4 header that follows (if applicable). Support for match criteria may depend on hardware or filter direction, as described below. Nokia recommends not configuring a filter in a direction or on hardware where a match criterion is not supported as this may lead to unwanted behavior. Some match criteria may be grouped in match lists and may be auto-generated based on router configuration. See [Filter Policy Advanced Topics](#) for more information.

Basic Layer 3 match criteria:

- **dscp** — Match for the specified DSCP value against the Differentiated Services Code Point/Traffic Class field in the IPv4 or IPv6 packet header.
- **src-ip/dst-ip** — Match for the specified source/destination IPv4/IPv6 address prefix against the source/destination IPv4/IPv6 address field in the IPv4/IPv6 packet header. The operator can optionally configure a mask to be used in a match.
- **flow-label** — Match for the specified flow label against the Flow label field in IPv6 packets. The operator can optionally configure a mask to be used in a match. This operation is supported on ingress filters.

Fragmentation match criteria:

- **fragment** — Enable fragmentation support in the filter policy match. For IPv4, match against the MF bit or Fragment Offset field to determine whether the packet is a fragment. For IPv6 for the 7750 SR and 7950 XRS, match against the Next Header Field for Fragment Extension Header value to determine whether the packet is a fragment. Up to six extension headers are matched against to find the Fragmentation Extension Header.

Additionally, IPv6 filters support mating against initial fragment using *first-only* or non-initial fragment *non-first-only*.

IPv4 match fragment criteria are supported on both ingress and egress. IPv6 match fragment criteria are supported on ingress only.

IPv4 options match criteria:

- **ip-option** — Matches the specified option value in the first option of the IPv4 packet. Operator can optionally configure a mask to be used in a match.
- **option-present** — Matches the presence of IP options in the IPv4 packet. Padding and EOOL are also considered as IP options. Up to six IP options are matched against.
- **multiple-option** — Matches the presence of multiple IP options in the IPv4 packet.
- **src-route-option** — Matches the presence of IP Option 3 or 9 (Loose or Strict Source Route) in the first three IP options of the IPv4 packet. A packet will also match this rule if the packet has more than three IP options.

IPv6 next-header match criteria: (see the Upper-layer protocol match next-header description below):

- **ah-ext-header** — Matches for the presence of the Authentication Header extension header in the IPv6 packet. This match criterion is supported on ingress only. Up to six extension headers are matched against.
- **esp-ext-header** — Matches for the presence of the Encapsulating Security Payload extension header in the IPv6 packet. This match criterion is supported on ingress only. Up to six extension headers are matched against.
- **hop-by-hop-opt** — Matches for the presence of Hop-by-hop options extension header in the IPv6 packet. This match criterion is supported on ingress only. Up to six extension headers are matched against.
- **routing-type0** — Matches for the presence of Routing extension header type 0 in the IPv6 packet. This match criterion is supported on ingress only. Up to six extension headers are matched against.

Upper-layer protocol match criteria:

- **next-header** — Matches the specified upper-layer protocol (such as, TCP, UDP, IGMPv6) against the Next Header field of the IPv6 packet header. “\*” can be used to specify TCP or UDP upper-layer protocol match (Logical OR). Next-header matching allows also matching on presence of a subset of IPv6 extension headers. See CLI section for details on which extension header match is supported.
- **protocol** — Matches the specified protocol against the Protocol field in the IPv4 packet header (for example, TCP, UDP, IGMP) of the outer IPv4. “\*” can be used to specify TCP or UDP upper-layer protocol match (Logical OR).
- **icmp-code** — Matches the specified value against the Code field of the ICMP/ICMPv6 header of the packet. This match is supported only for entries that also define protocol/next-header match for “ICMP”/“ICMPv6” protocol.
- **icmp-type** — Matches the specified value against the Type field of the ICMP/ICMPv6 header of the packet. This match is supported only for entries that also define protocol/next-header match for “ICMP”/“ICMPv6” protocol.
- **src-port/dst-port/port** — Matches the specified port value, port list, or port range against the Source Port Number/Destination Port Number of the UDP/TCP/SCTP packet header. An option to match either source or destination (Logical OR) using a single filter policy entry is supported by using a directionless “port” command. Source/destination match is supported only for entries that also define protocol/next-header match for “TCP”, “UDP”, “SCTP”, or “TCP or UDP” protocols. A non-initial fragment will never match an entry with non-zero port criteria specified.
- **tcp-ack/tcp-syn** — Matches the TCP ACK/TCP SYNC flag presence/absence in the TCP header of the packet. This match is supported only for entries that also define protocol/next-header match for “TCP” protocol.

Operational note for fragmented traffic — IP and IPv6 filters defined to match TCP, UDP, ICMP, or SCTP criteria (such as **src-port**, **dst-port**, **port**, **tcp-ack**, **tcp-syn**, **icmp-type**, and **icmp-code**) with values of zero or false will also match non-first fragment packets if other match criteria within the same filter entry are also met. Non-initial fragment packets do not contain a UDP, TCP, ICMP or SCTP header.

### 4.1.1.3 MAC Filter Policy Entry Match Criteria

The following list describes the MAC match criteria supported by SR OS or switches for all types of MAC filters (normal, isid, and vid). The criteria are evaluated against the Ethernet header of the Ethernet frame. Support for a match criteria may depend on H/W and/or filter direction as per below description. Match criterion is blocked if it is not supported by a specified frame-type or MAC filter sub-type. Nokia recommends not configuring a filter in a direction or on hardware where a match condition is not supported as this may lead to unwanted behavior.

- **frame-type** — The filter searches to match a specific type of frame format. For example, configuring frame-type ethernet\_II will match only Ethernet-II frames.
- **src-mac** — The filter searches to match source MAC address frames. Operator can optionally configure a mask to be used in a match.
- **dst-mac** — The filter searches to match destination MAC address frames. Operator can optionally configure a mask to be used in a match.
- **dot1p** — The filter searches to match 802.1p frames. The operator can optionally configure a mask to be used in a match.
- **etype** — The filter searches to match Ethernet II frames. The Ethernet type field is a two-byte field used to identify the protocol carried by the Ethernet frame.
- **ssap** — The filter searches to match frames with a source access point on the network node designated in the source field of the packet. Operator can optionally configure a mask to be used in a match.
- **dsap** — The filter searches to match frames with a destination access point on the network node designated in the destination field of the packet. Operator can optionally configure a mask to be used in a match.
- **snap-oui** — The filter searches to match frames with the specified three-byte OUI field.
- **snap-pid** — The filter searches to match frames with the specified two-byte protocol ID that follows the three-byte OUI field.
- **isid** — The filter searches to match for the matching Ethernet frames with the 24-bit ISID value from the PBB I-TAG. This match criterion is mutually exclusive of all the other match criteria under a specific MAC filter policy and is applicable to MAC filters of type **isid** only. The resulting MAC filter can only be applied on a BVPLS SAP or PW in the egress direction.
- **inner-tag/outer-tag** — The filter searches to match Ethernet frames with the non-service delimiting tags, as described in the [VID MAC Filters](#) section. This match criterion is mutually exclusive with all other match criteria under a specific MAC filter policy and is applicable to MAC filters of type **vid** only.

#### 4.1.1.4 IP Exception Filters

An NGE node supports IPv4 exception filters. For information on IP exception filters, refer to the “Router Encryption Exceptions using ACLs” section in the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR Layer 3 Services Guide: IES and VPRN*.



---

### 4.1.1.5 Filter Policy Actions

The actions are supported by ACL filter policies:

- **drop** — Allows operators to deny traffic to ingress or egress the system.
  - **IPv4 packet-length and IPv6 payload-length conditional drop** — Traffic can be dropped based on IPv4 packet length or IPv6 payload length by specifying a packet length or payload length value or range within the drop filter action (the IPv6 payload length field does not account for the size of the fixed IP header, which is 40 bytes).

This filter action is supported on ingress IPv4 and IPv6 filter policies only, if the filter is configured on an egress interface the **packet-length** or **payload-length** match condition is always true.

The additional match condition is part of action evaluation, such as, after the packet is determined to match the entry based on other configured match criteria.

Packets that match a filter policy entry match criteria and the **drop packet-length-value** or **payload-length-value** are dropped. Packets that match only the filter policy entry match criteria and do not match the **drop packet-length-value** or **drop payload-length-value** are forwarded with no further match in following filter entries.

Interaction with cflowd, log and mirror: The filter entry supports cflowd and log regardless of the outcome of the rate limit while forwarded packets only are mirrored.

- **IPv4 TTL and IPv6 hop limit conditional drop** — Traffic can be dropped based on IPv4 TTL or IPv6 hop limit by specifying a ttl or hop limit value or range within the **rate-limit** filter action.

This filter action is supported on ingress IPv4 and IPv6 filter policies only. If the filter is configured on an egress interface the packet-length or payload-length match condition is always true.

The additional match condition is part of action evaluation, such as, after the packet is determined to match the entry based on other match criteria configured.

Packets that match filter policy entry match criteria and the **drop ttl** or **hop-limit-value** are dropped. Packets that match only the filter policy entry match criteria and do not match the **drop ttl** or **hop-limit-value** are forwarded with no further match in following filter entries.

Interaction with cflowd, log and mirror: The filter entry supports cflowd and log regardless of the outcome of the rate limit while forwarded packets only are mirrored.

- **drop-extracted-traffic** — Traffic extracted to the CPM can be dropped using ingress IPv4 and IPv6 filter policies based on filter match criteria. Any IP traffic extracted to the CPM is subject to this filter action, including routing protocols, snooped traffic, and TTL expired traffic.

Packets that match the filter entry match criteria and extracted to the CPM are dropped. Packets that match only the filter entry match criteria and are not extracted to the CPM are forwarded with no further match in the subsequent filter entries.

Cflowd, log, mirror, and statistics apply to all traffic matching the filter entry, regardless of **drop** or **forward** action.

- **forward** — Allows operators to permit traffic to ingress or egress the system and be subject to regular processing.
- **rate-limit** — This action allows operators to rate-limit traffic matching a filter entry match criteria using IPv4, IPv6, or MAC filter policies.

If multiple interfaces (including LAG interfaces) use the same **rate-limit** filter policy on different FPs, then the system allocates a rate limiter resource for each FP; an independent rate limit applies to each FP.

If multiple interfaces (including LAG interfaces) use the same **rate-limit** filter policy on the same FP, then the system allocates a single rate limiter resource to the FP; a common aggregate rate limit is applied to those interfaces.

Note that traffic extracted to the CPM is not rate limited by an ingress **rate-limit** filter policy while any traffic generated by the router can be rate limited by an egress **rate-limit** filter policy.

Interaction with cflowd, log and mirror: **rate-limit** filter policy entries can coexist with cflowd, log, and mirror regardless of the outcome of the rate limit.

Interaction with QoS: Packets matching an ingress **rate-limit** filter policy entry bypass ingress QoS queuing or policing, and only the filter rate limit policer is applied. Packets matching an egress **rate-limit** filter policy bypass egress QoS policing, normal egress QoS queuing still applies.

- **IPv4 packet-length and IPv6 payload-length rate limit** — Traffic can be rate limited based on the IPv4 packet length and IPv6 payload length by specifying a packet-length value or payload-length value or range within the **rate-limit** filter action. The IPv6 payload-length field does not account for the size of the fixed IP header, which is 40 bytes.

This filter action is supported on ingress IPv4 and IPv6 filter policies only and cannot be configured on egress access or network interfaces.

This additional rate-limit condition is part of the filter entry action evaluation, it is not part of the filter entry match evaluation. It is performed after the packet is determined to match the entry based on the configured filter entry match criteria.

Packets that match a filter policy's entry match criteria and the **rate-limit packet-length-value** or **rate-limit payload-length-value** are rate limited. Packets that match only the filter policy's entry match criteria and do not match the **rate-limit packet-length-value** or **rate-limit payload-length-value** are forwarded with no further match in subsequent filter entries.

Cflowd, logging, and mirroring apply to all traffic matching the ACL entry regardless of the outcome of the rate limiter and regardless of the **packet-length-value** or **payload-length-value**.

- **IPv4 TTL and IPv6 hop-limit rate limit** — Traffic can be rate limited based on the IPv4 TTL or IPv6 hop-limit by specifying a TTL or hop limit value or range within the **rate-limit** filter action using ingress IPv4 or IPv6 filter policies.

This additional rate limit condition is part of the filter entry action evaluation. It is not part of the filter entry match evaluation. The evaluation is performed after the packet is determined to match the entry based on the configured filter entry match criteria.

The additional match condition is part of action evaluation (for example, after the packet is determined to match the entry based on other match criteria configured). Packets that match a filter policy entry match criteria and the **rate-limit ttl** or **hop-limit** value are rate limited. Packets that match only the filter policy entry match criteria and do not match the **rate-limit ttl** or **hop-limit** value are forwarded with no further matching in the subsequent filter entries.

Cflowd, logging, and mirroring apply to all traffic matching the ACL entry regardless of the outcome of the **rate-limit** value and the **ttl-value** or **hop-limit-value**.

- **forward** “Policy-based Routing/Forwarding (PBR/PBF) action” — Allows operators to permit ingress traffic but change the regular routing/forwarding that a packet would be a subject to. The PBR/PBF is applicable to unicast traffic only. The following PBR/PBF actions are supported (See CLI section for command details):

- **egress-pbr** — Enabling **egress-pbr** activates a PBR action on egress, while disabling **egress-pbr** activates a PBR action on ingress (default).

The following subset of the PBR actions (defined as follows) can be activated on egress: **redirect-policy**, **next-hop router**, and **esi**.

Egress PBR is supported in IPv4 and IPv6 filter policies for ESM only. Unicast traffic that is subject to slow-path processing on ingress (for example, IPv4 packets with options or IPv6 packets with hop-by-hop extension header) will not match egress PBR entries. Filter logging, cflowd, and mirror source are mutually exclusive of configuring a filter entry with an egress PBR action. Configuring **pbr-down-action-override**, if supported with a specific PBR ingress action type, is also supported when the action

is an egress PBR action. Processing defined by **pbr-down-action-override** does not apply if the action is deployed in the wrong direction. If a packet matches a filter PBR entry and the entry is not activated for the direction in which the filter is deployed, **action forward** is executed. Egress PBR cannot be enabled in system filters.

- **esi** — Forwards the incoming traffic using VXLAN tunnel resolved using EVPN MP BGP control plane to the first service chain function identified by ESI (Layer 2) or ESI/SF-IP (Layer 3). Supported with VPLS (Layer 2) and IES/VRPN (Layer 3) services. If the service function forwarding cannot be resolved, traffic matches an entry and **action forward** is executed.

For VPLS, no cross-service PBF is supported; that is, the filter specifying ESI PBF entry must be deployed in the VPLS service where BGP EVPN control plane resolution takes place as configured for a specific ESI PBF action. The functionality is supported in filter policies deployed on ingress VPLS interfaces. BUM traffic that matches a filter entry with ESI PBF will be unicast forwarded to the VTEP:VNI resolved through PBF forwarding.

For IES/VRPN, the outgoing R-VPLS interface can be in any VRPN service. The outgoing interface and VRPN service for BGP EVPN control plane resolution must again be configured as part of ESI PBR entry configuration. The functionality is supported in filter policies deployed on ingress IES/VRPN interfaces and in filter policies deployed on ingress and egress for ESM subscribers. Only unicast traffic is subject to ESI PBR; any other traffic matching a filter entry with Layer 3 ESI action will be subject to **action forward**.

When deployed in unsupported direction, traffic matching a filter policy ESI PBR/PBF entry will be subject to **action forward**.

- **lsp** — Forwards the incoming traffic onto the specified LSP. Supports RSVP-TE LSPs (type **static** or **dynamic** only), MPLS-TP LSPs, or SR-TE LSPs. Supported for ingress IPv4/IPv6 filter policies and only deployed on IES SAPs or network interfaces. If the configured LSP is down, traffic matches the entry and **action forward** is executed.
- **next-hop** — Changes the IP destination address used in routing from the address in the packet to the address configured in this PBR action. The operator can configure whether the next-hop IP address must be direct (local subnet only) or indirect (any IP). This functionality is supported for ingress IPv4/IPv6 filter policies only, and is deployed on Layer 3 interfaces. If the configured next-hop is not reachable, traffic is dropped and a “ICMP destination unreachable” message is sent. If the indirect keyword is not specified but the IP address is a remote IP address, traffic will be dropped.
  - **interface** — Forwards the incoming traffic onto the specified IPv4 interface. Supported for ingress IPv4 filter policies in global routing table instance. If the configured interface is down or not of the supported type, traffic is dropped.

- **redirect-policy** — Implements PBR **next-hop** or PBR **next-hop router** action with ability to select and prioritize multiple redirect targets and monitor the specified redirect targets so PBR action can be changed if the selected destination goes down. Supported for ingress IPv4 and IPv6 filter policies deployed on Layer 3 interfaces only. See section [Redirect Policies](#) for further details.
- **remark dscp** — Allows an operator to remark the DiffServ Code Points of packets matching filter policy entry criteria. Packets are remarked regardless of QoS-based in-/out-of- profile classification and QoS-based DSCP remarking is overridden. DSCP remarking is supported both as a main action and as an extended action. As a main action, this functionality applies to IPv4 and IPv6 filter policies of any scope and can only be applied at ingress on either access or network interfaces of Layer 3 services only. As an extended action, this functionality applies to IPv4 and IPv6 filter policies of any scope and can be applied at ingress on either access or network interfaces of Layer 3 services, or at egress on Layer 3 subscriber interfaces. The functionality requires IOM3 or above.
- **router** — Changes the routing instance a packet is routed in from the upcoming interface's instance to the routing instance specified in the PBR action (supports both GRT and VPRN redirect). It is supported for ingress IPv4/IPv6 filter policies deployed on Layer 3 interfaces. The action can be combined with the **next-hop** action specifying direct/indirect IPv4/IPv6 next hop. Packets are dropped if they cannot be routed in the configured routing instance. For further details, see section "Traffic Leaking to GRT" in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN*.
- **sap** — Forwards the incoming traffic onto the specified VPLS SAP. Supported for ingress IPv4/IPv6 and MAC filter policies deployed in VPLS service. The SAP that the traffic is to egress on must be in the same VPLS service as the incoming interface. If the configured SAP is down, traffic is dropped.
- **sdp** — Forwards the incoming traffic onto the specified VPLS SDP. Supported for ingress IPv4/IPv6 and MAC filter policies deployed in VPLS service. The SDP that the traffic is to egress on must be in the same VPLS service as the incoming interface. If the configured SDP is down, traffic is dropped.
- **vprn-target** — Redirects the incoming traffic in a similar manner to combined next-hop and LSP redirection actions, but with greater control and slightly different behavior. This action is supported for both IPv4 and IPv6 filter policies and is applicable on ingress of access interfaces of IES/ VPRN services. See [Filter Policy Advanced Topics](#) for further details.
- forward "isa action" — ISA processing actions allow operator to permit ingress traffic and send it for ISA processing as per specified ISA action. The following ISA actions are supported (see CLI section for command details):

- **gtp-local-breakout** — Forwards matching traffic to NAT instead of being GTP tunneled to the mobile operator's PGW or GGSN. The action applies to GTP-subscriber-hosts. If filter is deployed on other entities, **action forward** is applied. Supported for IPv4 ingress filter policies only. If ISAs performing NAT are down, traffic is dropped.
- **nat** — Forwards matching traffic for NAT. Supported for IPv4/IPv6 filter policies for Layer 3 services in GRT or VPRN. If ISAs performing NAT are down, traffic is dropped. (see CLI for options)
- **reassemble** — Forwards matching packets to the reassembly function. Supported for IPv4 ingress filter policies only. If ISAs performing reassemble are down, traffic is dropped.
- **tcp-mss-adjust** — Forwards matching packets (TCP Syn) to an ISA BB Group for MSS adjustment. In addition to the IP filter, the operator also needs to configure the **mss-adjust-group** command under the Layer 3 service to specify the *bb-group-id* and the new *segment-size*.
- **http-redirect** — Implements the HTTP redirect captive portal. HTTP GET is forwarded to CPM card for captive portal processing by router. See the [HTTP-redirect \(Captive Portal\)](#) section for more information.

In addition to the above actions:

- An operator can select a **default-action** for a filter policy. The default action is executed on packets subjected to an active filter when none of the filter's active entries matches the packet. By default, filter policies have default action set to drop but operator can select a default action to be forward instead.
- An operator can override default action applied to packets matching a PBR/PBF entry when the PBR/PBF target is down using **pbr-down-action-override**. Supported options are to drop the packet, forward the packet, or apply the same action as configured for the filter policy's **default-action**. The override is supported for the following PBR/PBF actions. For the last three actions, the override is supported whether in redundancy mode or not.
  - **forward esi** (Layer 2 or Layer 3)
  - **forward sap**
  - **forward sdp**
  - **forward next-hop [indirect] router**

[Table 46](#) defines default behavior for packets matching a PBR/PBF filter entry when a target is down.

**Table 46** Default behavior when a PBR/PBF target is down

| PBR/PBF action         | Default behavior when down |
|------------------------|----------------------------|
| forward esi (any type) | Forward                    |

**Table 46** Default behavior when a PBR/PBF target is down (Continued)

| PBR/PBF action              | Default behavior when down                                                            |
|-----------------------------|---------------------------------------------------------------------------------------|
| forward lsp                 | Forward                                                                               |
| forward next-hop (any type) | Drop                                                                                  |
| forward redirect-policy     | Forward when redirect policy is shutdown                                              |
| forward redirect-policy     | Forward when destination tests are enabled and the best destination is not reachable  |
| forward redirect-policy     | Drop when destination tests are not enabled and the best destination is not reachable |
| forward sap                 | Drop                                                                                  |
| forward sdp                 | Drop                                                                                  |
| forward router              | Drop                                                                                  |
| forward vprn-target         | Forward                                                                               |

#### 4.1.1.6 Viewing Filter Policy Actions

A number of parameters determine the behavior of a packet after it has been matched to a defined criterion or set of criteria:

- the action configured by the user
- the context in which a filter policy is applied. For example, applying a filter policy in an unsupported context can result in simply forwarding the packet rather than applying the configured action.
- external factors, such as the reachability (according to a given test criteria) of a target

Because of this, SR OS provides the following commands that enable the user to capture this context globally and identify how a packet will be handled by the system:

- **show>filter>ip**
- **show>filter>ipv6**
- **show>filter>mac**

This section describes the key information displayed as part of the output for the **show** commands listed above, and explains how to interpret it.

---

From a configuration point of view, the **show** command output displays the main action (primary and secondary), as well as the extended action.

The “PBR Target Status” field shows the basic information that the system has of the target based on simple verification methods. This information is only shown for the filter entries which are configured in redundancy mode (that is, with both primary and secondary main actions configured), and for ESI redirections. Specifically, the target status in the case of redundancy depends on several factors; for example, on a match in the routing table for next-hop redirects, or on VXLAN tunnel resolution for ESI redirects.

The “Downloaded Action” field specifically describes the action that the system will perform on the packets that match the criterion (or criteria). This typically depends on the context in which the filter has been applied (whether it is supported or not), but in the case of redundancy, it also depends on the target status. For example, the downloaded action will be the secondary main action when the target associated to the primary action is down. In the nominal (for example, non-failure condition) case the “Downloaded Action” will reflect the behavior a packet will be subject to. However, in transient cases (for example, in the case of a failure) it may not be able to capture what will effectively happen to the packet.

The output also displays relevant information such as the default action when the target is down (see [Table 46](#)) as well as the overridden default action when **pbr-down-action-override** has been configured.

There are situations where, collectively, this information does not capture what will effectively happen to the packet throughout the system. To that end, the **effective-action** keyword of the **show>filter>[ip | ipv6 | mac]** commands enables advanced checks to be performed and accurate packet fates to be displayed.

The criteria for determining when a target is down. While there is little ambiguity on that aspect when the target is local to the system performing the steering action, ambiguity is much more prominent when the target is distant. Therefore, because the use of **effective-action** triggers advanced tests, a discrepancy is introduced compared to the action when **effective-action** keyword is not used. This will, for example, be the case for redundant actions.



---

### 4.1.1.7 Filter Policy Statistics

Filter policies support per-entry, packet/byte match statistics. The cumulative matched packet/Byte counters are available per ingress and per egress direction. Every packet arriving on an interface/service/subscriber using a filter policy increments ingress or egress (as applicable) matched packet/Byte count for a filter entry the packet matches (if any) on the line card the packet ingresses/egresses. For each policy, the counters for all entries are collected from all line cards, summarized and made available to an operator.

Starting with SR OS Release 11.0R4, filter policies applied on access interfaces are downloaded only when active and only to line cards that have interfaces associated with those filter policies. If a filter policy is not downloaded to any line card, the statistics show 0. If a filter policy is being removed from any of the line cards the policy is currently downloaded to (as result of association change or when a filter becomes inactive), the statistics for the filter are reset to 0. Downloading a filter policy to a new line card keeps incrementing existing statistics.

Starting with SR OS Release 13.0R4, filter policies support bulk requests of CPM cache for policy interface-created entries. The cache is periodically refreshed through a background collection of counters from hardware. The counters are also refreshed when the ACL entry corresponding to the cache entry has statistics read from hardware through any direct-read from hardware mechanism. If a cache entry represents an entry for an ACL filter policy not downloaded to any line cards, the cache returns values of 0. If a cache entry represents an ACL filter entry that was removed from a line card since the previous refresh, the current refresh will reload the cache with the most recent values from hardware. The cache has to be rebuilt on a High Availability (HA) switchover, accordingly initial statistics requests after an HA switchover may require reads from hardware.

Operational notes:

- Two consecutive bulk requests for one entry will return the same values if the cache has not been refreshed between the two requests. The refresh interval is platform/release dependent. Contact your Nokia representative for more information.
- The cache is currently used only for Open Flow statistics retrieval. See [Hybrid OpenFlow Switch](#) for more details.
- Conditional action match criteria filter entries for **ttl**, **hop-limit**, **packet-length**, and **payload-length** support logging and statistics when the condition is met, allowing visibility of filter matched and action executed. If the condition is not met, packets are not logged and statistics against the entry are not incremented.

### 4.1.1.8 Filter Policy Logging

SR OS supports logging of the information from the packets that match a specific filter policy. Logging is configurable per filter policy entry by specifying preconfigured filter log (**config>filter>log**). A filter log can be applied to ACL filters and CPM hardware filters. Operators can configure multiple filter logs and specify: memory allocated to a filter log destination, syslog ID for filter log destination, filter logging summarization, and wrap-around behavior.

Notes related to filter log summarization:

- The implementation of the feature applies to filter logs with destination syslog.
- Summarization logging is the collection and summarization of log messages for one specific log ID within a period of time.
- The summarization interval is 100 seconds.
- Upon activation of a summary, a mini-table with src/dst-address and count is created for each type (IPv4/IPv6/MAC).
- Every received log packet (due to filter match) is examined for source or destination address.
- If the log packet (source/destination address) matches a source/destination address entry in the mini-table, from a packet received previously, the summary counter of the matching address is incremented.
- If source or destination address of the log messages does not match an entry already present in the table, the source/destination address is stored in a free entry in the mini-table.
- In case the mini-table has no more free entries, only total counter is incremented.
- At expiry of the summarization interval, the mini-table for each type is flushed to the syslog destination.

Operational note:

- Conditional action match criteria filter entries for **ttl**, **hop-limit**, **packet-length**, and **payload-length** support logging and statistics when the condition is met, allowing visibility of filter matched and action executed. If the condition is not met, packets are not logged and statistics against the entry are not incremented.

### 4.1.1.9 Filter Policy cflowd Sampling

Filter policies can be used to control how cflowd sampling is performed on an IP interface. If an IP interface has cflowd sampling enabled, an operator can exclude some flows for interface sampling by configuring filter policy rules that match the flows and by disabling interface sampling as part of the filter policy entry configurations (**interface-disable-sample**). If an IP interface has cflowd sampling disabled, an operator can enable cflowd sampling on a subset of flows by configuring filter policy rules that match the flows and by enabling cflowd sampling as part of the filter policy entry configurations (**filter-sample**).

The above cflowd filter sampling behavior is exclusively driven by match criteria. The sampling logic applies regardless of whether an action was executed (including evaluation of conditional action match criteria, for example, **packet-length** or **ttl**).

### 4.1.1.10 Filter Policy Management

#### 4.1.1.10.1 Modifying Existing Filter Policy

There are several ways to modify an existing filter policy. A filter policy can be modified through configuration change or can have entries populated through dynamic, policy-controlled dynamic interfaces; for example, RADIUS, OpenFlow, flowspec, or Gx. Although in general, SR OS ensures filter resources exist before a filter can be modified, because of the dynamic nature of the policy-controlled interfaces, a configuration that was accepted may not be applied in H/W due to lack of resources. When that happens, an error is raised.

A filter policy can be modified directly—by changing/adding/deleting the existing entry in that filter policy—or indirectly. Examples of indirect change to filter policy include changing embedded filter entry this policy embeds (see the [Embedded Filters](#) section) or changing redirect policy this filter policy uses.

Finally, a filter policy deployed on a specific interface can be changed by changing the policy the interface is associated with.

All of the above changes can be done in service. A filter policy that is associated with service/interface cannot be deleted unless all associations are removed first.

For a large (complex) filter policy change, it may take a few seconds to load and initiate the filter policy configuration. Filter policy changes are downloaded to line cards immediately; therefore, operators should use filter policy copy or transactional CLI to ensure partial policy change is not activated.

---

#### 4.1.1.10.2 Filter Policy Copy and Renumbering

To assist operators in filter policy management, SR OS supports entry copy and entry renumbering operations.

Filter **copy** allows operators to perform bulk operations on filter policies by copying one filter's entries to another filter. Either all entries or a specified entry of the source filter can be selected for copy. When entries are copied, entry order is preserved unless destination filter's entry ID is selected (applicable to single-entry copy). The filter copy allows overwrite of the existing entries in the destination filter by specifying "overwrite" option during the copy command. Filter copy can be used, for example, when creating new policies from existing policies or when modifying an existing filter policy (an existing source policy is copied to a new destination policy, the new destination policy is modified, then the new destination policy is copied back to the source policy with overwrite specified).

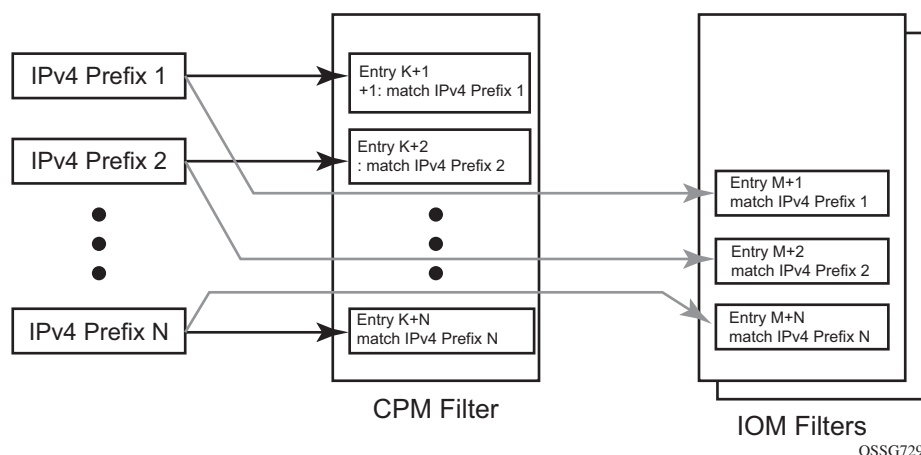
Entry renumbering allows operators to change relative order of a filter policy entry by changing the entry ID. Entry renumbering can also be used to move two entries closer together or further apart, thereby creating additional entry space for new entries.

### 4.1.2 Filter Policy Advanced Topics

#### 4.1.2.1 Match List for Filter Policies

[Figure 26](#) shows an approach to implement logical OR on a list of matching criteria (IPv4 address prefixes in this example) in one or more filter policies prior to introduction of match list.

**Figure 26 IOM/CPM Filter Policy Using Individual Address Prefixes**



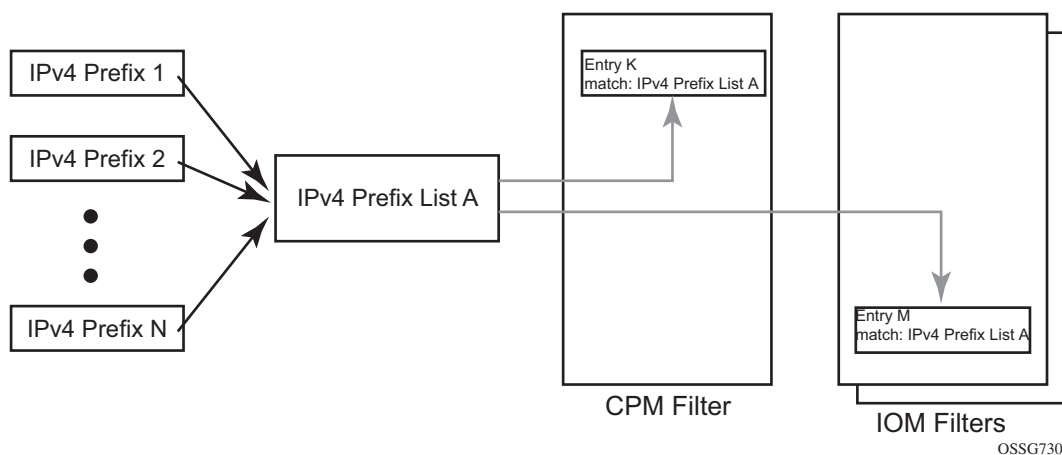
OSSG729

An operator has to create one entry for each address prefix to execute a common action. Each entry defines a match on a unique address prefix from the list plus any other additional match criteria and the common action. If the same set of address prefixes needs to be used in another IOM/line card, or CPM filter policy, an operator again needs to create one entry for each address prefix from the list in those filter policies. Same procedure applies (not shown above) if another action needs to be performed on the list of the addresses within the same filter policy (when, for example, specifying different additional match criteria). This process can introduce large operational overhead, especially when a list contains many elements or needs to be reused multiple times across one or more filter policies.

Match lists for CPM and IOM/FP filter policies eliminate the preceding operational complexity by simplifying the IOM/FP and CPM filter policy management on a list of match criteria. Instead of defining multiple filter entries in any specific filter, an operator can now group the same types of matching criteria into a single filter match list and use that list as a match criterion value, thus requiring only a single filter policy entry per each unique action. The same match list can be used in one or more IOM/line card filter policies as well as CPM filter policies.

The match lists further simplify management and deployment of the policy changes. A change in a match-list content is automatically propagated across all policies employing that list in their match criteria, therefore, only a single configuration change is required to trigger policy changes when a list is used by multiple entries in one or more filter policies.

Figure 27 depicts how the IOM/CPM filter policy changes with a filter match list usage (using IPv4 address prefix list in this example).

**Figure 27 IOM/CPM Filter Policy Using an Address Prefix Match List**

The hardware resource usage does not change whether filter match lists are used or whether operator creates multiple entries (each per one element of the list): however, a careful consideration must be given to how the lists are used to ensure only needed match permutations are created in a filter policy entry (especially when other matching criteria that are also lists or ranges are specified in the same entry). The system verifies that a new list element, for example, an IP address prefix, cannot be added to a specific list or a list cannot be used by a new filter policy unless resources exist in hardware to implement the required filter policy (ies) that reference that list. If that is not the case, addition of a new element to the list or use of the list by another policy will fail.

Some use cases, like those driven by dynamic policy changes, may result in acceptance of filter policy configuration changes that cannot be programmed in hardware because of the resource exhaustion. If that is the case, when attempting to program a filter entry that uses match lists, the operation will fail, the entry will not be programmed, and a notification of that failure will be provided to an operator.

Refer to the Software Release Notes for information about objects that can be grouped into a filter match list for FP and CPM filter policies.

#### 4.1.2.1.1 Apply-Path — Auto-Generation of Prefix List Entries

Using the filter **match-list apply-path** commands, the router supports the auto-generation of IPv4 and IPv6 prefix list entries from the router configuration for BGP peers configured in the GRT or VPRN. This capability is often required to simplify the management of CPM or line card filters.

Using the filter **match-list apply-path** capability, the operator can:

- specify one or more regex expression matches against the SR OS configuration per list
- specify wildcard matches by specifying the regex wildcard match expression (“.”)
- mix auto-generated entries with statically configured entries within a match list

Additional rules followed when using apply-path are:

- Operational and administrative states of a specific router configuration are ignored when auto-generating address prefixes.
- Duplicates are not removed when populated by different auto-generation matches and static configuration.
- A configuration will fail if auto-generation of an address prefix would result in filter policy resource exhaustion on a filter entry, system, or line card level.

#### 4.1.2.2 Embedded Filters

When a large number of standard filter policies are configured in a system, a set of policies will often contain one or more common blocks of entries that define, for example, system-wide and/or service-wide security rules. Before introduction of the embedded filters, such common rules would have to be configured separately in each exclusive/template policy.

To simplify management of such common rules across multiple filter policies, the operator can use embedded filter policies. An embedded filter policy is a special type of a filter policy that cannot be deployed directly but instead is used to define a common filter policy rules that are then included in (embedded into) other filter policies in the system. Thanks to embedding, a common set of rules can now be defined and changed in a single place but deployed across multiple filter policies.

The following main rules apply when embedding an embedded filter policy:

1. An operator can explicitly define an offset at which to embed a specific embedded filter into a specific embedding filter—the embedded filter entry number X becomes an entry (X + offset) in the embedding filter.
2. An exclusive/template filter policy may embed multiple embedded filter policies as long as the embedded entries do not overlap.
3. A single embedded filter policy may be embedded in many exclusive/template filter policies.

4. When embedding an embedded filter, an operator may want to change or deactivate an embedded filter policy entry in the embedding filter, allowing for customization of the common embedded filter policy rules by the embedding filter. This can be achieved by either defining an entry in the embedding filter that will match ahead of the embedded filter entry or by overwriting the embedded filter entry in the embedding filter.


For example: If embedded filter 99 has entry 20 that drops packets that match IP source address **src\_address**, and filter 200 embeds filter 99 at offset 100, then to *deactivate* the embedded entry 20, an operator could define an entry 120 (embedded entry number 20 + offset 100) in filter policy 200 that has the same match criteria and has either:

- no action defined (this will deactivate the embedded entry and allow continued evaluation of filter policy 200)
  - **action forward** defined (packets will match the new entry and will be forwarded instead of dropped, evaluation of filter policy 200 will stop)
5. Any embedded policy rule edits are automatically applied to all filter policies that embed that embedded filter policy.
  6. The system verifies whether system and h/w resources exist when a new embedded filter policy is created, changed or embedded. If resources are not available, the configuration is rejected. In rare cases, filter policy resource check may pass but filter policy can still fail to load due to a resource exhaustion on a line card (for example when other filter policy entries are dynamically configured by applications like RADIUS in parallel). If that is the case, the embedded filter policy configured will be deactivated (configuration will be changed from **activate** to **inactivate**).
  7. An embedded filter is never embedded partially into an exclusive/template filter; that is, resources must exist to embed all embedded filter entries in a specific exclusive/template filter. Although a partial embedding into a single filter will not take place, an embedded filter may be embedded only in a subset of embedding filters (only those where there are sufficient resources available).

[Figure 28](#) shows implementation of embedded filter policy using IPv4 ACL filter policy example with an embedded filter 10 being used to define common filter rules that are then embedded into filter 1 and 20 (with filter 20 overwriting rule at offset 50).



**Figure 28 Embedded Filter Policy**



**Note:** Embedded filter policies are supported for line card IP(v4) and IPv6 filter policies only.

### 4.1.2.3 System-level IPv4/IPv6 Line Card Filter Policy

A system filter policy allows the definition of a common set of policy rules that can then be activated within other exclusive/template filters. IPv4/IPv6 system filter policies supports all IPv4/IPv6 filter policy match rules and actions respectively but system policy entries cannot be the sources of mirroring.

System filter policy cannot be used directly; the active system policy is deployed by activating it within any IPv4 or IPv6 exclusive/template filter policy (chaining the system policy and a specific interface policy). When an IPv4/IPv6 filter policy is chained to the active IPv4/IPv6 system filter, system filter rules are evaluated first before any rules of the chaining filter are evaluated (i.e. chaining filter's rules are only matched against if no system filter match took place).

A system filter policy is intended mainly for system-level blacklisting rules, therefore it is recommended to use system policies with drop/forward actions. Other actions like, for example, PBR actions, or redirect to ISAs should not be used unless the system filter policy is activated only in filters used by services that support such action. The “nat” action is not supported and should not be configured. Failure to observe these restrictions can lead to unwanted behavior as system filter actions are not verified against the services the chaining filters are deployed for.

System filter policies can be populated using CLI/SNMP/Netconf management interfaces and Openflow policy interface. System filter policy entries cannot be populated using flowspec, RADIUS, or Gx.

System filter policy scale is identical to a corresponding IPv4 or IPv6 filter policy scale. System filter policy consumes single set of H/W resources on each line card as soon as it is activated, regardless of how many IPv4/IPv6 filters chain to that system policy. This optimizes resource allocation when multiple filter policies activate a specific system policy.

An example (IPv4) configuration is shown below:

```
*A:vm1>config>filter#
Configure system-policy
 ip-filter 1 create
 scope system
 entry 5 create
 match protocol *
 fragment true
 exit
 action drop
 exit
exit
Activate it
system-filter
 ip 1
exit
Use it in another filter:
ip-filter 10 create
 chain-to-system-filter
 filter-name "test-name"
 embed-filter open-flow "test" offset 100
exit
exit
```

#### 4.1.2.4 Primary and Secondary Filter Policy Action for PBR/PBF Redundancy

In some deployments, operators may want to specify a backup PBR/PBF target if the primary target is down. SR OS allows the configuration of a primary action (**config>filter>{ip-filter | ipv6-filter | mac-filter}>entry>action**) and a secondary action (**config>filter>{ip-filter | ipv6-filter | mac-filter}>entry>action secondary**) as part of a single filter policy entry. The secondary action can only be configured if the primary action is configured.

For Layer 2 PBF redundancy, the operator can configure the following redundancy options:

- **action forward sap AND action secondary forward sap**

- **action forward sdp AND action secondary forward sdp**
- **action forward sap AND action secondary forward sdp**
- **action forward sdp AND action secondary forward sap**

For Layer 3 PBR redundancy, an operator can configure any of the following actions as a primary action and any (either same or different than primary) of the following as a secondary action. Furthermore, none of the parameters need to be the same between primary and secondary actions. Although the following commands refer to IPv4 in the *ip-address* parameter, they also apply to IPv6.

- **forward next-hop *ip-address* router *router-instance***
- **forward next-hop *ip-address* router *service-name* *service-name***
- **forward next-hop indirect *ip-address* router *router-instance***
- **forward next-hop indirect *ip-address* router *service-name* *service-name***

When primary and secondary actions are configured, PBR/PBF uses the primary action if its target is operationally up, or it uses the secondary action if the primary PBR/PBF target is operationally down. If both targets are down, the default action when the target is down (see [Table 46](#)), as per the primary action, is used, unless **pbr-down-action-override** is configured.

When PBR/PBF redundancy is configured, the operator can use sticky destination functionality for a redundant filter entry. When sticky destination is configured (**config>filter>{ip-filter | ipv6-filter | mac-filter}>entry>sticky-dest**), the functionality mimics that of sticky destination configured for redirect policies. To force a switchover from the secondary to the primary action when sticky destination is enabled and secondary action is selected, the operator can use the **tools>perform>filter>{ip-filter | ipv6-filter | mac-filter}>entry>activate-primary-action** command. Sticky destination can be configured even if no secondary action is configured.

The control plane monitors whether primary and secondary actions can be performed and programs forwarding filter policy to use either the primary or secondary action as required. More generally, the state of PBR/PBF targets is monitored in the following situations:

- when a secondary action is configured
- when sticky destination is configured
- when a **pbr-down-action-override** is configured

The **show>filter>{ip-filter | ipv6-filter | mac-filter} [entry]** command displays which redundant action is activated or downgraded, including when both PBR/PBF targets are down. The following example shows partial output of the command as applicable for PBF redundancy.

```
*A:vsim-200001# show filter ip 10 entry 1000
...
Primary Action : Forward (SAP) <-details of (primary) action
 Next Hop : 1/1/1
 Service Id : Not configured
 PBR Target Status : Does not exist
Secondary Action : Forward (SAP) <-details of secondary action
 Next Hop : 1/1/2
 Service Id : Not configured
 PBR Target Status : Does not exist
PBR Down Action : Forward (pbr-down-action-override) <- PBR down behavior
Downloaded Action : None <- currently downloaded action
Dest. Stickiness : 1000 Hold Remain : 0 <-
 sticky dest details
```

#### 4.1.2.5 Extended Action for Performing Two Actions at a Time

In certain deployment scenarios, for example to realize service function chaining, operators may want to perform a second action in addition to a traffic steering action. SR OS allows this behavior by configuring an extended action for a main action. This functionality is supported for Layer 3 traffic steering (that is, PBR) and more specifically for the following main actions:

- **forward esi** (Layer 3 version)
- **forward lsp**
- **forward next-hop [indirect] [router]**
- **forward next-hop interface**
- **forward redirect-policy**
- **forward router**
- **forward vprn-target**

Furthermore, the capability to specify an extended action is also supported in the case of PBR redundancy, therefore for the following action:

- **forward next-hop [indirect] router**

The supported extended action is:

- **remark dscp** *dscp-name*

The extended action is available in the following contexts: **config>filter>ip-filter>entry>action>extended-action** and **config>filter>ipv6-filter>entry>action>extended-action**.

If the status of the target of the main action is tracked, which is the case, amongst others, for PBR/PBF redundancy, the extended action listed above will not be performed when the PBR target is down. Moreover, a filter policy containing an entry with the extended action **remark dscp** will be blocked in the following cases: if applied on ingress with the **egress-pbr flag** set, if applied on egress without the **egress-pbr flag** set. The latter case includes actions that are not supported on egress (and for which **egress-pbr** cannot be set).

#### 4.1.2.6 Advanced VPRN Redirection

The **vprn-target** action is a resilient redirection capability which combines both data-path and control plane lookups to achieve the desired redirection. It allows for the following redirection models:

- redirection towards the default PE while selecting a specific LSP to use
- redirection towards an alternative PE while selecting or not a specific LSP to use. If a specific LSP is not selected, then the system will automatically select one based on the BGP next-hop tunnel resolution mechanism
- all of the above within any VPRN

When configuring this action, the user must specify the target BGP next-hop (**bgp-nh**) towards which the redirection should occur, as well as the routing context (**router**) in which the necessary lookups will be performed (to derive the service label).

The target BGP next-hop can be configured with any label allocation method (label per VRF, label per next-hop, label per prefix). These methods entail different forwarding behaviors; however, the steering node is not aware of the configuration of the target node. If the user does not specify an advertised route prefix (**adv-prefix**), the steering node will assume that label per VRF is used by the target node and will select the service label accordingly. If the target node is not operating according to the label per VRF method, the user must specify an appropriate route prefix for which a service label is advertised by the target node, keeping in mind the resulting forwarding behavior at the target node of the redirected packet. This specification will instruct the steering node to use that specific service label.

Be aware that the system will perform an exact match between the specified *ip-address/mask* (or *ipv6-address/prefix-length*) and the advertised route.

The user can specify an LSP (RSVP-TE, MPLS-TP, or SR-TE LSP) to use towards the BGP next-hop. If no LSP is specified, the system will automatically select one the same way it would have done when normally forwarding a packet towards the BGP next-hop.



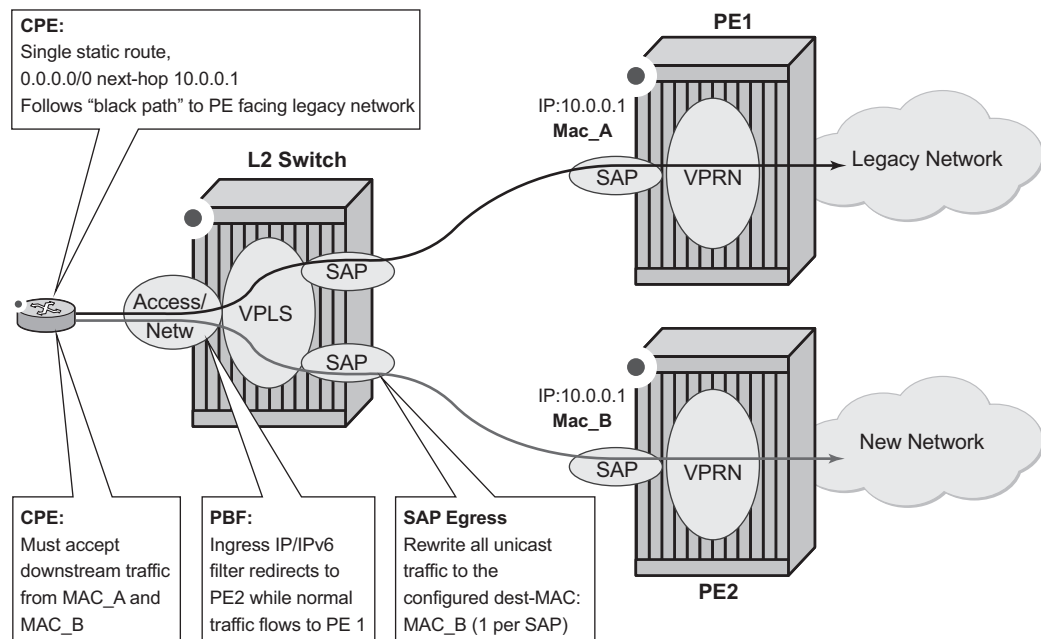
**Note:** While the system only performs the redirection when the traffic is effectively able to reach the target BGP next-hop, it does not verify whether the redirected packets will effectively reach their destination after that.

This action is resilient in that it tracks events affecting the redirection at the service level and reacts to those events. As such, the system will perform the redirection as long as it can reach the target BGP next-hop using the proper service label. If the redirection cannot be performed (for example, if no LSP is available, the peer is down, or there is no more specific labeled route), the system will revert to normal forwarding. This can be overridden and configured to drop. A maximum of 8k of unique (3-tuple {**bgp-nh**, **router**, **adv-prefix**}) redirection targets can be tracked.

#### 4.1.2.7 Destination MAC Rewrite When Deploying Policy-Based Forwarding

For Layer 2 Policy-Based Forwarding (PBF) redirect actions, a far-end router may discard redirected packets when the PBF changes the destination IP interface the packet arrives on. This happens when a far-end IP interface uses a different MAC address than the IP interface reachable via normal forwarding (for example, one of the routers does not support a configurable MAC address per IP interface). To avoid the discards, operators can deploy egress destination MAC rewrite functionality for VPLS SAPs (**config>service>vpls>sap>egress>dest-mac-rewrite**). [Figure 29](#) shows a sample deployment.

**Figure 29 Layer 2 Policy-Based Forwarding (PBF) redirect action**



0920

When enabled, all unicast packets have their destination MAC rewritten to operator-configured value on an Layer 2 switch VPLS SAP. Multicast and broadcast packets are unaffected. The feature:

- Is supported for regular and split-horizon group Ethernet SAPs in a regular VPLS Service
- Is expected to be deployed on a SAP that faces far-end IP interface (either a SAP that is the target of PBF action, as shown in [Figure 29](#), or a VPLS SAP of a downstream Layer 2 switch that is connected to a far-end router—not shown).
- Applies to any unicast egress traffic including LI and mirror.

Restrictions:

- Is mutually exclusive with SAP MAC ingress and egress loopback feature: **tools perform service-id service-id loopback eth sap sap-id {ingress | egress} mac-swap ieee-address**

### 4.1.2.8 Network-port VPRN Filter Policy

Network-port Layer 3 service-aware filter feature allows operators to deploy VPRN service aware ingress filtering on network ports. A single ingress filter of **scope template** can each be defined for IPv4 and for IPv6 against a VPRN service. The filter applies to all unicast traffic arriving on auto-bind and explicit-spoke network interfaces for that service. The network interface can be either Inter-AS, or Intra-AS. The filter does not apply to traffic arriving on access interfaces (SAP, spoke-sdp, network-ingress (CsC, rVPLS, eVPN).

The same filter can be used on access interfaces of the specific VPRN, can embed other filters (including OpenFlow), can be chained to a system filter, and can be used by other Layer 2 or Layer 3 services.

The filter is deployed on all line cards (chassis network mode D is required). There are no limitations related to filter match/action criteria or embedding. The filter is programmed on line cards against ILM entries for this service. All label-types are supported. If an ILM entry has a filter index programmed, that filter is used when the ILM is used in packet forwarding; otherwise, no filter is used on the service traffic.

Restrictions:

- Network port Layer 3 service-aware filters do not support flowspec and LI (cannot use filter inside LI infrastructure nor have LI sources within the VPRN filter).

### 4.1.2.9 ISID MAC Filters

ISID filters are a type of MAC filters that allows filtering based on the ISID values rather than Layer 2 criteria used by MAC filters of type "**normal**" or "**vid**". ISID filters can be deployed on iVPLS PBB SAPs and ePipe PBB SAPs in the following scenarios:

The MMRP usage of the MRP policy ensures automatically that traffic using Group BMAC is not flooded between domains. However, there could be small transitory periods when traffic originated from PBB BEB with unicast BMAC destination may be flooded in the BVPLS context as unknown unicast in the BVPLS context for both iVPLS and PBB Epipe. To restrict distribution of this traffic for local PBB services, ISID filters can be deployed. The MAC filter configured with ISID match criterion can be applied to the same interconnect endpoints (BVPLS SAP or PW) as the MRP policy to restrict the egress transmission of any type of frames that contains a local ISID. The ISID filters will be applied as required on a per B-SAP or B-PW basis, just in the egress direction.



The ISID match criteria are exclusive with any other criteria under **mac-filter**. A new **mac-filter** type attribute is defined to control the use of ISID match criteria and must be set to ISID to allow the use of ISID match criteria.

#### 4.1.2.10 VID MAC Filters

VID filters are a type of MAC filters that extend the capability of current Ethernet ports with null or default SAP tag configuration to match and take action on VID tags. Service delimiting tags (for example, QinQ 1/1/1:10.20 or dot1q 1/1/1:10, where outer tag 10 and inner tags 20 are service delimiting) allow fine granularity control of frame operations based on the VID tag. Service delimiting tags are exact match and are stripped from the frame as shown in [Figure 30](#). Exact match or service delimiting tags do not require VID filters. VID filters can only be used to match on frame tags that are after the service delimiting tags.

With VID filters, operators can choose to match VID tags for up to two tags on ingress, egress, or both.

- The outer tag is the first tag in the packet that is carried transparently through the service.
- The inner tag is the second tag in the packet that is carried transparently through the service.

VID filters add the capability to perform VID value filter policies on default tags (1/1/1:\*, or 1/1/1:x.\*, or 1/1/1:\*.0) or null tags (1/1/1, 1/1/1:0, or 1/1/1:x.0). The matching is based on the port configuration and the SAP configuration.

At ingress, the system looks for the two outer-most tags in the frame. If present, any service delimiting tags are removed and not visible to VID MAC filtering. For example:

- 1/1/1:x.y SAP has no tag left for VID MAC filter to match on (outer-tag and inner-tag = 0)
- 1/1/1:x.\* SAP has potentially one tag in the \* position for VID MAC filter to match on
- SAP such as 1/1/1, 1/1/1:\*, or 1/1/1:\*. \* can have as many as two tags for VID MAC filter to match on
- For the remaining tags, the left (outer-most) tag is what is used as the outer tag in the MAC VID filter. The following tag is used as the inner tag in the filter. If any of these positions do not have tags, a value of 0 is used in the filter. At egress, the VID MAC filter is applied to the frame prior to adding the additional service tags.

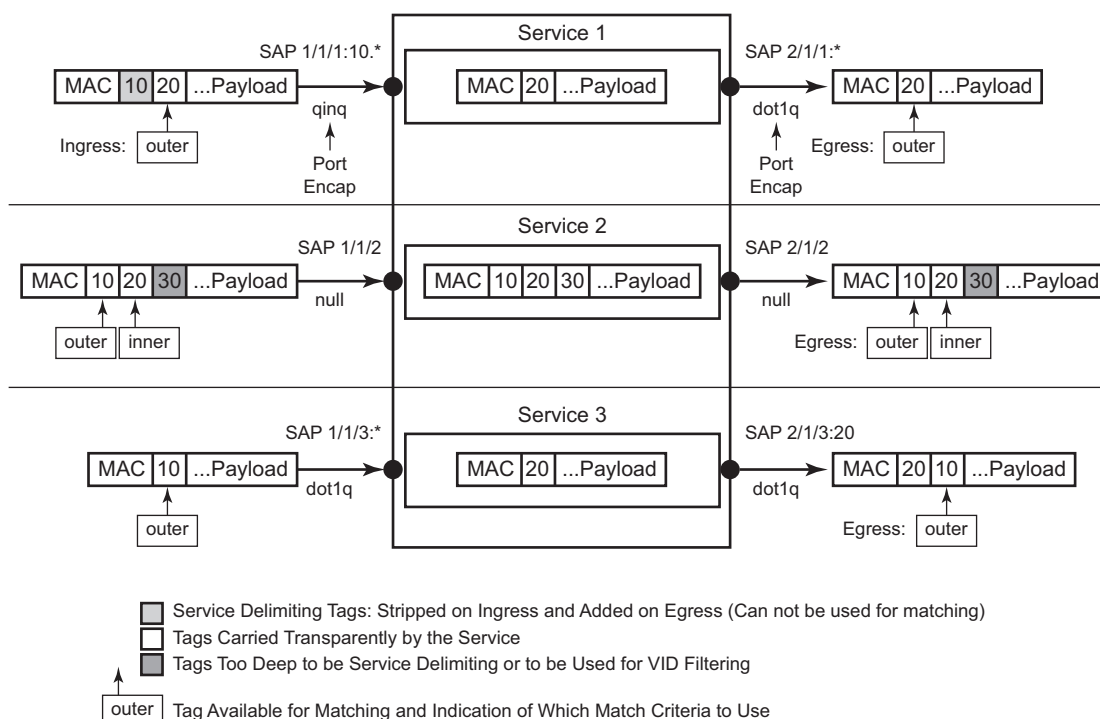
In the industry, the QinQ tags are often referred to as the C-VID (customer VID) and S-VID (service VID). The terms outer tag and inner tag allow flexibility without having to refer to C-TAG and S-TAG explicitly. The position of inner and outer tags is relative to the port configuration and SAP configuration. Matching of tags is allowed for up to the first two tags on a frame because service delimiting tags may be 0, 1, or 2 tags.

The meaning of inner and outer has been designed to be consistent for egress and ingress when the number of non-service delimiting tags is consistent. Service 1 in [Figure 30](#) shows a conversion from QinQ to a single dot1q example where there is one non-service delimiting tag on ingress and egress. Service 2 shows a symmetric example with two non-service delimiting tags (plus and additional tag for illustration) to two non-service delimiting tags on egress. Service 3 shows a single non-service delimiting tag on ingress and two tags with one non-service delimiting tag on ingress and egress.

SAP-ingress QoS setting allows for MAC-criteria type VID, which uses the VID filter matching capabilities of QoS and VID Filters (see the *7450 ESS*, *7750 SR*, *7950 XRS*, and *VSR Quality of Service Guide*).

A VID filter entry can also be used as a debug or lawful intercept mirror source entry.

**Figure 30** VID Filtering Examples



OSSG735

---

VID filters are available on Ethernet SAPs for Epipe, VPLS, or I-VPLS including eth-tunnel and eth-ring services.

#### 4.1.2.10.1 Arbitrary Bit Matching of VID Filters

In addition to matching an exact value, a VID filter mask allows masking any set of bits. The masking operation is  $((\text{value and vid-mask}) = (\text{tag and vid-mask}))$ . For example: A value of 6 and a mask of 7 would match all VIDs with the lower 3 bits set to 6. VID filters allow explicit matching of VIDs and matching of any bit pattern within the VID tag.


When using VID filters on SAPs, only VID filters are allowed on this SAP. Filters of type normal and ISID are not allowed.

An additional check for the "0" VID tag may be required when using certain wild card operations. For example, frames with no tags on null encapsulated ports will match a value of 0 in outer tag and inner tag because there are no tags in the frame for matching. If a zero tag is possible but not wanted, it can be explicitly filtered using exact match on "0" prior to testing other bits for "0".

**config>system>ethernet>new-qinq-untagged-sap** is a special QinQ function for single tagged QinQ frames with a null second tag. Using this in combination with VID filters is not recommended. The outer tag is the only tag available for filtering on egress for frames arriving from MPLS SDPs or from PBB services, even though additional tags may be carried transparently.

### 4.1.2.10.2 Port Group Configuration Example

**Figure 31 Port Groups**



OSSG734

Figure 31 shows a customer use example where some VLANs are prevented from ingressing or egressing certain ports. In the example, port A sap 1/1/1:1.\* would have a filter as shown below while port A sap 1/1/1:2.\* would not.:

```
mac-filter 4 create
 default-action forward
 type vid
 entry 1 create
 match frame-type ethernet_II
 outer-tag 30 4095
 exit
 action drop
 exit
exit
```

### 4.1.2.11 IP Exception Filters

IP exception filters scan all outbound traffic entering a Network Group Encryption (NGE) domain and allow packets that match the exception filter criteria to transit the NGE domain unencrypted.

The VSR supports IPv4 exception filters. For information on IP exception filters, refer to the “Router Encryption Exceptions using ACLs” section in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN*.

The most basic IP exception filter policy must have the following:

- an exception filter policy ID
- scope, either exclusive or template
- at least one filter entry with a specified matching criteria

### 4.1.2.12 Redirect Policies

SR OS-based routers support configuring of IPv4 and IPv6 redirect policies. Redirect policies allow specifying multiple redirect target destinations and defining status check test methods used to validate the ability for a destination to receive redirected traffic. This destination monitoring allows routers to react to target destination failures. To specify an IPv4 redirect policy, define all destinations to be IPv4. To specify an IPv6 redirect policy, define all destinations to be IPv6. IPv4 redirect policies can only be deployed in IPv4 filter policies. IPv6 redirect policy can only be deployed in IPv6 filter policies.

Redirect policies support the following destination tests:

- **ping test** – with configurable interval, drop-count, and time-out
- **url-test** – with configurable URL, interval, drop-count, timeout, and configurable action (disable destination, lower or raise priority) based upon return error code
- **snmp-test** – with configurable OID and Community strings, interval, drop-count, timeout, and configurable action (disable destination, lower or raise priority) based upon SNMP return value.
- **unicast-rt-test** – unicast routing reachability, supported only when router instance is configured for a specific redirect policy. The test yields true if the route to the specified destination exists in RTM for the configured router instance.

Each destination is assigned an initial or base priority describing this destination’s relative importance within the policy. The destination with the highest priority value is selected as most-preferred destination and programmed on line cards in filter policies using this redirect policy as an action. Only destinations that are not disabled by the programmed test (if configured) are considered when selecting the most-preferred destination.

In some deployments, it may not be necessary to switch from a currently active, most-preferred redirect-policy destination when a new more-preferred destination becomes available. To support such deployments, operators can enable the sticky destination functionality (**config>filter>redirect-policy>sticky-dest**). When enabled, the currently active destination remains active unless it goes down or an operator forces the switch using the **tools>perform>filter>redirect-policy>activate-best-dest** command. An optional sticky destination *hold-time-up* is available to delay programming the sticky destination in the redirect policy (transition from **action forward** to PBR action to the most-preferred destination). When the timer is enabled, the first destination that comes up will not be programmed and instead the timer is started. Once the timer expires, the most-preferred destination at that time will be programmed (which may be a different destination from the one that started the timer). Note the following:

- When the manual switchover to most-preferred destination is executed as described above, the hold-time-up is stopped.
- When the timer value is changed, the new value takes immediate effect and the timer is restarted with the new value (or expired if **no-hold-time-up** is configured).



**Note:** The **unicast-rt-test** command will fail when performed in the context of a VPRN routing instance when the destination is routable only through **grt-leak** functionality. **ping-test** is recommended in such cases.

Feature restrictions:

- Redirect policy is supported for ingress IPv4 and IPv6 filter policies only.
- SNMP and URL tests are not supported for IPv6.
- Different platforms support different scale for redirect policies. Contact your local Nokia representative to ensure the planned deployment does not exceed recommended scale.

#### 4.1.2.12.1 Router Instance Support for Redirect Policies

There are two modes of deploying redirect policies on VPRN interfaces. The functionality supported depends on the configuration of the **redirect-policy** router option with **config>filter>redirect-policy>router**:

- Redirect policy with router option enabled (recommended):

- When a PBR destination is up, the PBR lookup is performed in the redirect policy's configured routing instance. When that instance differs from the incoming interface where the filter policy using the specific redirect policy is deployed, the PBR action is equivalent to forward next-hop router filter policy action.
- When all PBR destinations are down (or a hardware does not support action router), **action forward** is programmed and the PBR lookup is performed in the routing instance of the incoming interface where the filter policy using the specific redirect policy is deployed.
- Any destination tests configured are executed in the routing context specified by the redirect policy.
- Changing router configuration for a redirect policy brings all destinations with a test configured down. The destinations are brought up once the test confirms reachability based on the new redirect policy router configuration.
- Redirect policy without router option disabled (**no router**) or with router options not supported (legacy):
  - When a PBR destination is up, the PBR lookup is performed in the routing instance of the incoming interface where the filter policy using the specific redirect policy is deployed.
  - When all PBR destinations are down, **action forward** is programmed and the PBR lookup is performed in the routing instance of the incoming interface where the filter policy using the specific redirect policy is deployed.
  - Any destination tests configured are always executed in the "Base" router instance regardless of the router instance of the incoming interface where the filter policy using the specific redirect policy is deployed.

Restrictions:

- Only **unicast-rt-test** and **ping-test** are supported when the router option is enabled.

#### 4.1.2.13 HTTP-redirect (Captive Portal)

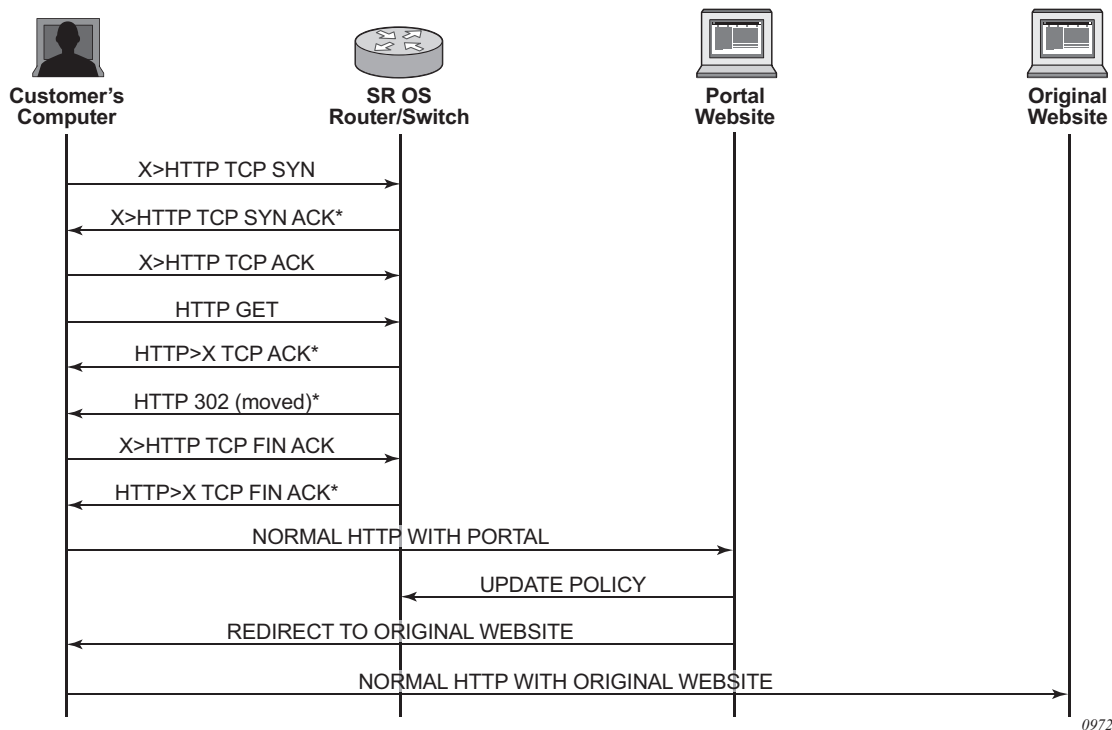
Web redirection policies can be configured on SR OSs/switches. The http redirection action can block a customer's request from an intended recipient and force the customer to connect to the service's portal server. 255 unique entries with **http-redirect** are allowed.

#### 4.1.2.13.1 Traffic Flow

The following example provides a brief scenario of a customer connection with web redirection.

1. The customer gets an IP address using DHCP (if the customer is trying to set a static IP he will be blocked by the anti-spoofing filter).
2. The customer tries to connect to a website.
3. The router intercepts the HTTP GET request and blocks it from the network
4. The router then sends the customer an HTTP 302 (service temporarily unavailable/moved). The target URL should then include the customer's IP and MAC addresses as part of the portal's URL.
5. The customer's web browser will then close the original connection and open a new connection to the web portal.
6. The web portal updates the ACL (directly or through SSC) to remove the redirection policy.
7. The customer connects to the original site.

**Figure 32 Web Redirect Traffic Flow**



Starred entries (\*) are items the router performs masquerading as the destination, regardless of the destination IP address or type of service.



The following displays information that can optionally be added as variables in the portal URL (http-redirect url):

- \$IP – The customer's IP address.
- \$MAC – The customer's MAC address.
- \$URL – The original requested URL.
- \$SAP – The customer's SAP.
- \$SUB – The customer's subscriber identification string".
- \$CID — A string that represents the circuit-id or interface-id of the subscriber host (hexadecimal format).
- \$RID — A string that represents the remote-id of the subscriber host (hexadecimal format).
- \$SAPDESC – A configurable string that represents the configured SAP description.

The subscriber identification string is available only when used with subscriber management. Refer to the subscriber management section of the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide* and the *7450 ESS, 7750 SR, 7950 XRS, and VSR Router Configuration Guide*.

Since most web sites are accessed using the domain name the router allows either DNS queries or responds to DNS with the portal's IP address.

#### 4.1.2.14 Filter Policies and Dynamic Policy-Driven Interfaces

In addition to configuration interfaces like CLI/SNMP, filter policies can be created and modified by dynamic policy-driven interfaces, such as BGP flowspec, OpenFlow, RADIUS, or XMPP-Python.

For BGP flowspec, routes are learned by a routing instance, and the system auto-creates an embedded filter to contain the rules derived from these routes. The maximum number of rules in the embedded filter of each routing instance can be controlled through configuration. The embedded filter containing the flowspec rules of a routing instance can be inserted into any configured exclusive or template-scope IPv4/IPv6 filter, and the embedding is activated if:

- the filter is applied to the ingress context of an IP interface that supports flowspec
- the IP interfaces to which the filter is applied all belong to the same routing instance, and that routing instance is the one associated with the flowspec routes

The insertion point of the flowspec rules in each embedding filter policy is controlled through offset configuration. For more information, see the BGP flowspec section of the *7450 ESS, 7750 SR, 7950 XRS, and VSR Unicast Routing Protocols Guide*.

For RADIUS, operator can assign filter policies to a subscriber, and populate filter policies used by the subscriber within a preconfigured block reserved for RADIUS filter entries. See the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide* and filter RADIUS-related commands for more details.

VSD filters are created dynamically via XMPP and managed via Python script so rules can be inserted into or removed from the correct VSD template or embedded filters. XMPP messages received by the 7750 SR are passed transparently to the Python module to generate the appropriate CLI. More information about VSD filter provisioning, automation, and Python scripting details are in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN*.

For OpenFlow, embedded filter infrastructure is used to inject OpenFlow rules into an existing filter policy. See [Hybrid OpenFlow Switch](#) for more details.

Policy-controlled auto-created filters are re-created on system reboot. Policy-controlled filter-entries are lost on system reboot and need to be reprogrammed.

#### 4.1.2.15 Filter Policy-based ESM Service Chaining

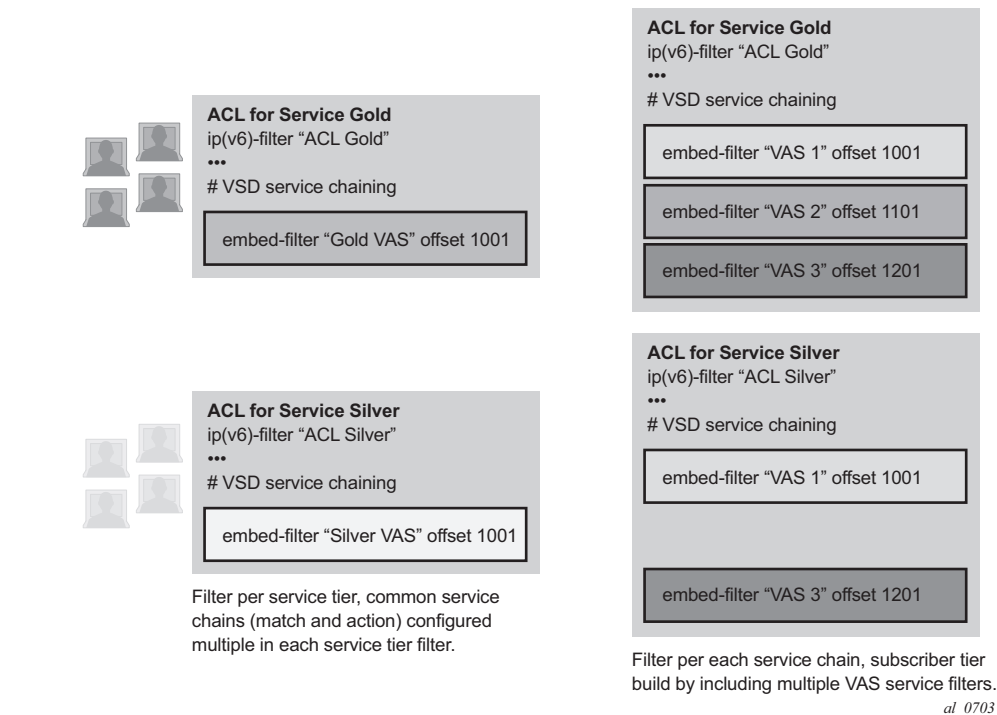
In some deployments, operators may select to redirect ESM subscribers to Value Added Services (VAS). Various deployment models can be used but often subscribers are assigned to a specific residential tier-of-service, which also defines the VAS available to subscribers of the specific tier. The subscribers are redirected to VAS based on tier-of-service rules but such an approach can be hard to manage when many VAS services/tiers of service are desired. Often the only way to identify a subscriber's traffic with a specific tier-of-service is to preallocate IP/IPv6 address pools to a specific service tier and use those addresses in VAS PBR match criteria. This creates an application-services to network infrastructure dependency that can be hard to overcome, especially if fast and flexible application service delivery is desired.

Filter policy-based ESM service chaining removes ESM VAS steering to network infrastructure inter-dependency. An operator can configure per tier of service or per individual VAS service upstream and downstream service chaining rules without a need to define subscriber or tier-of-service match conditions. [Figure 33](#) shows a possible ACL model (embedded filters are used for VAS service chaining rules).

On the left in [Figure 33](#), the per-tier-of-service ACL model is depicted. Each tier of service (Gold or Silver) has a dedicated embedded VAS filter (“Gold VAS”, “Silver VAS”) that contains all steering rules for all service chains applicable to the specific tier. Each VAS filter is then embedded by the ACL filter used by a specific tier. A subscriber is subject to VAS service chain rules based on the per-tier ACL assigned to that subscriber (for example, via RADIUS). If a new VAS rule needs to be added, an operator must program that rule in all applicable tiers. Upstream and downstream rules can be configured in a single filter (as shown) or can use dedicated ingress and egress filters.

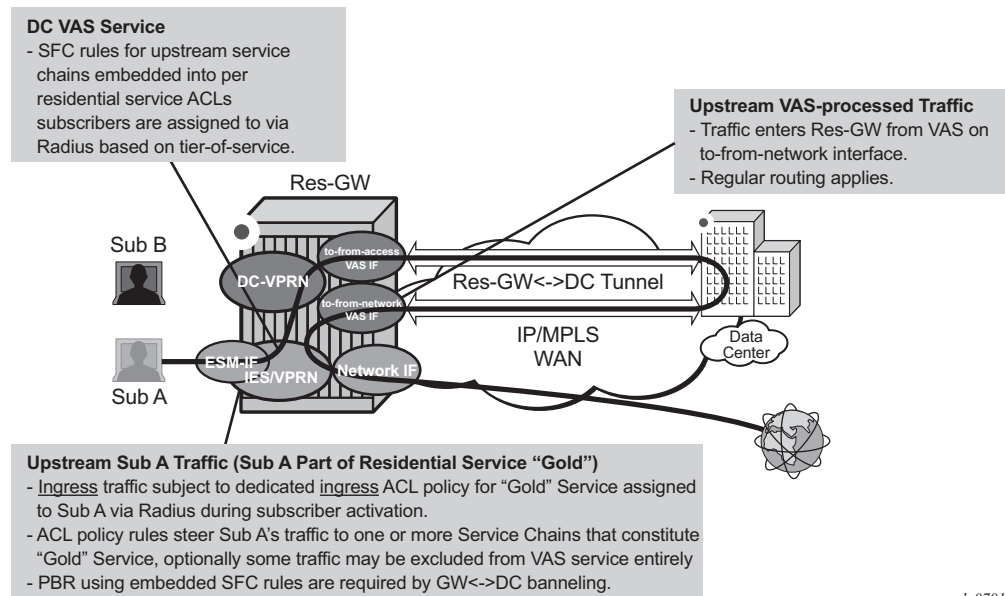
On the right in [Figure 33](#), the per-VAS-service ACL model is depicted. Each VAS has a dedicated embedded filter (“VAS 1”, “VAS 2”, “VAS 3”) that contains all steering rules for all service chains applicable to that VAS service. A tier of service is then created by embedding multiple VAS-specific filters: Gold: VAS 1, VAS 2, VAS 3; Silver: VAS 1 and VAS 3. A subscriber is subject to VAS service chain rules based on the per-tier ACL assigned to that subscriber. If a new VAS rule needs to be added, an operator needs to program that rule in a single VAS-specific filter only. Again, upstream and downstream rules can be configured in a single filter (as shown) or can use dedicated ingress and egress filters.

**Figure 33** ACL filter modeling for ESM Service Chaining



**Figure 34** shows upstream VAS service chaining steering using filter policies. Upstream subscriber traffic entering Res-GW is subject to the subscriber's ingress ACL filter assigned to that subscriber by a policy server. If the ACL contains VAS steering rules, the VAS-rule-matching subscriber traffic is steered for VAS processing over a dedicated to-from-access VAS interface in the same or a different routing instance. After the VAS processing, the upstream traffic can be returned to Res-GW by a to-from-network interface (shown in **Figure 34**) or can be injected to WAN to be routed toward the final destination (not shown).

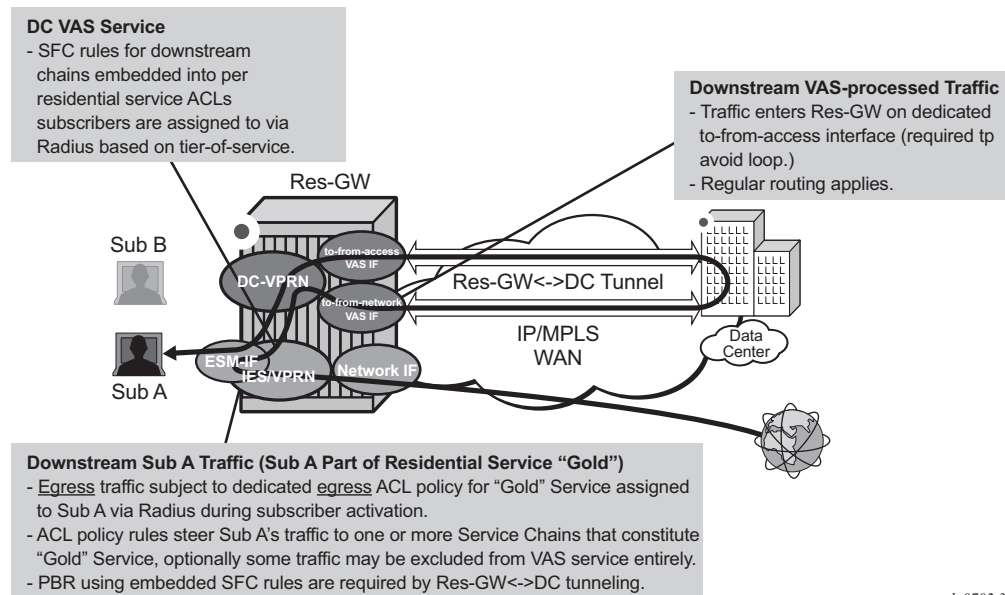
**Figure 34 Upstream ESM ACL-policy based service chaining**



al\_0701

**Figure 35** shows downstream VAS service chaining steering using filter policies. Downstream subscriber traffic entering Res-GW is forwarded to a subscriber-facing line card. On that card, the traffic is subject to the subscriber's egress ACL filter policy processing assigned to that subscriber by a policy server. If the ACL contains VAS steering rules, the VAS rule-matching subscriber's traffic is steered for VAS processing over a dedicated to-from-network VAS interface (in the same or a different routing instance). After the VAS processing, the downstream traffic must be returned to Res-GW via a "to-from-network" interface (shown in **Figure 35**) to ensure the traffic is not redirected to VAS again when the subscriber-facing line card processes that traffic.

**Figure 35 Downstream ESM ACL-policy based service chaining**



al\_0702.3

Ensuring the correct settings for the VAS interface type, for upstream and downstream traffic redirected to a VAS and returned after VAS processing, is critical for achieving loop-free network connectivity for VAS services. The available configuration options (**config>service>vprn>if>vas-if-type**, **config>service>ies>if>vas-if-type** and **config>router>if>vas-if-type**) are described below:

- deployments that use two separate interfaces for VAS connectivity (recommended, and required if local subscriber-to-subscriber VAS traffic support is required)
  - **to-from-access**
    - upstream traffic arriving from subscribers over access interfaces must be redirected to a VAS PBR target reachable over this interface for upstream VAS processing
    - downstream traffic destined for subscribers after VAS processing must arrive on this interface, so that the traffic is subject to regular routing but is not subject to Application Assurance diversion, nor to egress subscriber PBR
    - the interface must not be used for downstream pre-VAS traffic; otherwise, routing loops will occur
  - **to-from-network**

- downstream traffic destined for subscribers arriving from network interfaces must be redirected to a VAS PBR target reachable over this interface for downstream VAS processing
- upstream traffic after VAS processing, if returned to the router, must arrive on this interface so that regular routing can be applied
- deployments that use a single interface for VAS connectivity (optional, no local subscriber-to-subscriber VAS traffic support)
  - **to-from-both**
    - both upstream traffic arriving from access interfaces and downstream traffic arriving from the network are redirected to a PBR target reachable over this interface for upstream/downstream VAS processing
    - after VAS processing, traffic must arrive on this interface (optional for upstream), so that the traffic is subject to regular routing but is not subject to AA diversion, nor to egress subscriber PBR
    - the interface must be used for downstream pre-VAS traffic; otherwise, routing loops will occur

The ESM filter policy-based service chaining allows operators to do the following:

- Steer upstream and downstream traffic per-subscriber with full ACL-flow-defined granularity without the need to specify match conditions that identify subscriber or tier-of-service
- Steer both upstream and downstream traffic on a single Res-GW
- Flexibly assign subscribers to tier-of-service by changing the ACL filter policy a specific subscriber uses
- Flexibly add new services to a subscriber or tier-of-service by adding the subscriber-independent filter rules required to achieve steering
- Achieve isolation of VAS steering from other ACL functions like security through the use of embedded filters
- Deploy integrated Application Assurance (AA) as part of a VAS service chain—both upstream and downstream traffic is processed by AA before a VAS redirect
- Select whether to use IP-Src/IP-Dst address hash or IP-Src/IP-Dst address plus TCP/UDP port hash when LAG/ECMP connectivity to DC is used. Layer 4 inputs are not used in hash with IPv6 packets with extension headers present.

ESM filter policy-based traffic steering supports the following:

- IPv4 and IPv6 steering of unicast traffic using IPv4 and IPv6 ACLs
- **action forward redirect-policy** or **action forward next-hop router** for IP steering with TCAM-based load-balancing, fail-to-wire, and sticky destination

- **action forward esi sf-ip vas-interface router** for an integrated service chaining solution

Operational notes:

- Downstream traffic steered toward a VAS on the subscriber-facing IOM is reclassified (FC and profile) based on the subscriber egress QoS policy, and is queued toward the VAS based on the network egress QoS configuration. Packets sent toward VAS will not have DSCP remarked (since they are not yet forwarded to a subscriber). DSCP remarking based on subscriber's egress QoS profile will only apply to traffic ultimately forwarded to the subscriber (after VAS or not subject to VAS).
- If mirroring of subscriber traffic is configured using ACL entry/subscriber/SAP/port mirror, the mirroring will apply to traffic ultimately forwarded to subscriber (after VAS or not subject to VAS). Traffic that is being redirected to VAS cannot be mirrored using an ACL filter implementing PBR action (the same egress ACL filter entry being a mirror source and specifying egress PBR action is not supported).
- Use dedicated ingress and egress filter policies to prevent accidental match of an ingress PBR entry on egress, and vice-versa, that will result in forwarding/dropping of traffic matching the entry (based on the filter's default action configuration).

Restrictions:

- This feature is not supported with HSMDAs on subscriber ingress.
- This feature is not supported when the traffic is subject to non-AA ISA on Res-GW.
- Traffic that matches an egress filter entry with an egress PBR action cannot be mirrored, cannot be sampled using cflowd, and cannot be logged using filter logging while being redirected to VAS on a sub-facing line card.
- This feature is not supported with LAC/LNS ESM (PPPoE subscriber traffic encapsulated into or de-encapsulated from L2TP tunnels).
- This feature is not supported for system filter policies.

#### 4.1.2.16 Policy-Based Forwarding for Deep Packet Inspection in VPLS

The purpose policy-based forwarding is to capture traffic from a customer and perform a deep packet inspection (DPI) and forward traffic, if allowed, by the DPI.

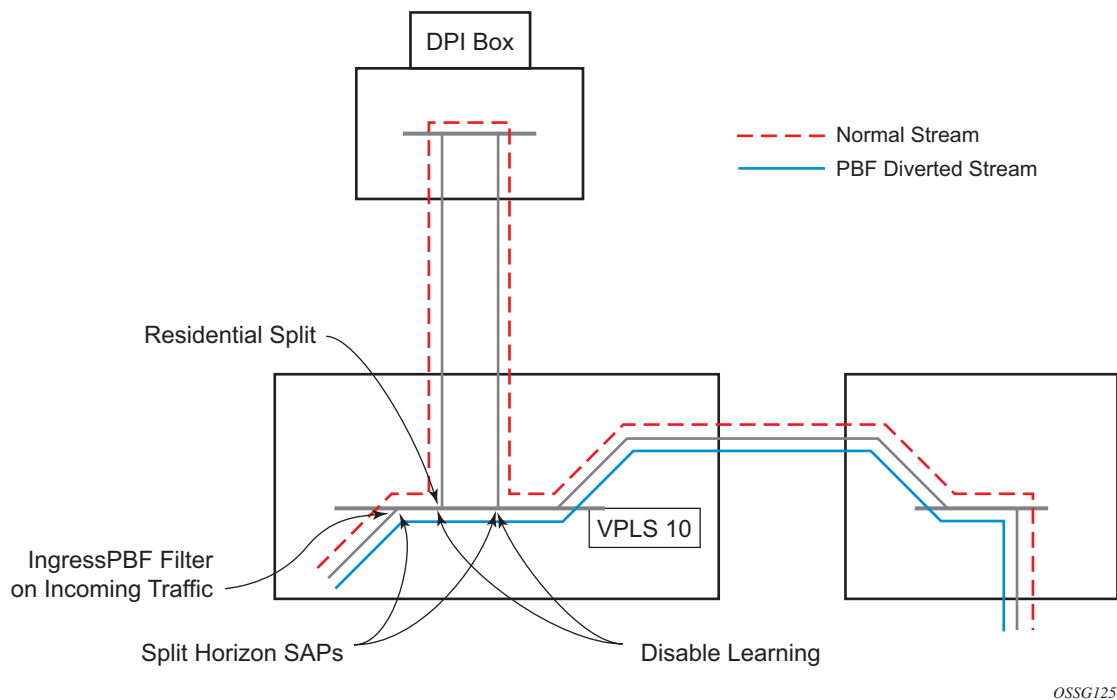
In the following example, the split horizon groups are used to prevent flooding of traffic. Traffic from customers enter at SAP 1/1/5:5. Due to the mac-filter 100 that is applied on ingress, all traffic with dot1p 07 marking will be forwarded to SAP 1/1/22:1, which is the DPI.

DPI performs packet inspection/modification and either drops the traffic or forwards the traffic back into the box through SAP 1/1/21:1. Traffic will then be sent to spoke-sdp 3:5.

SAP 1/1/23:5 is configured to see if the VPLS service is flooding all the traffic. If flooding is performed by the router then traffic would also be sent to SAP 1/1/23:5 (which it should not).

Figure 36 shows an example to configure policy-based forwarding for deep packet inspection on a VPLS service. For information about configuring services, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN*.

**Figure 36 Policy-Based Forwarding for Deep Packet Inspection**



The following displays a VPLS service configuration with DPI example:

```
*A:ALA-48>config>service# info

...
vpls 10 customer 1 create
```



```

 service-mtu 1400
 split-horizon-group "dpi" residential-group create
 exit
 split-horizon-group "split" create
 exit
 stp
 shutdown
 exit
 sap 1/1/21:1 split-horizon-group "split" create
 disable-learning
 static-mac 00:00:00:31:11:01 create
 exit
 sap 1/1/22:1 split-horizon-group "dpi" create
 disable-learning
 static-mac 00:00:00:31:12:01 create
 exit
 sap 1/1/23:5 create
 static-mac 00:00:00:31:13:05 create
 exit
 no shutdown
 exit
...

*A:ALA-48>config>service#

```

The following displays a MAC filter configuration example:

```

*A:ALA-48>config>filter# info

...
 mac-filter 100 create
 default-action forward
 entry 10 create
 match
 dot1p 7 7
 exit
 log 101
 action forward sap 1/1/22:1
 exit
 exit
...

*A:ALA-48>config>filter#

```

The following displays the MAC filter added to the VPLS service configuration:

```

*A:ALA-48>config>service# info

...
 vpls 10 customer 1 create
 service-mtu 1400
 split-horizon-group "dpi" residential-group create
 exit
 split-horizon-group "split" create
 exit
 stp
 shutdown
 exit

```

---

```
exit
sap 1/1/5:5 split-horizon-group "split" create
 ingress
 filter mac 100
 exit
 static-mac 00:00:00:31:15:05 create
exit
sap 1/1/21:1 split-horizon-group "split" create
 disable-learning
 static-mac 00:00:00:31:11:01 create
exit
sap 1/1/22:1 split-horizon-group "dpi" create
 disable-learning
 static-mac 00:00:00:31:12:01 create
exit
sap 1/1/23:5 create
 static-mac 00:00:00:31:13:05 create
exit
spoke-sdp 3:5 create
exit
no shutdown
exit
....

*A:ALA-48>config>service#
```

## 4.2 Configuring Filter Policies with CLI

This section provides information to configure filter policies using the command line interface.

### 4.2.1 Common Configuration Tasks

This section provides a brief overview of the tasks that must be performed for both IP and MAC filter configurations and provides the CLI commands.

#### 4.2.1.1 Creating an IPv4 Filter Policy

Configuring and applying filter policies is optional. Each filter policy must have the following:

- The filter type specified (IP)
- A filter policy ID
- A default action, either drop or forward
- Filter policy scope specified, either *exclusive* or *template*
- At least one filter entry with matching criteria specified
- Optionally, an existing filter policy can have a Filter Name assigned, that can then be used in CLI to reference that filter policy including assigning it to SAPs and/or network interfaces.

Use the following CLI syntax to create a template IPv4 filter policy:

**CLI Syntax:** `config>filter# ip-filter filter-id [create]  
                  description description-string  
                  scope {exclusive | template}  
                  default-action {drop | forward}`

##### 4.2.1.1.1 IPv4 Filter Entry

Within a filter policy, configure filter entries which contain criteria against which ingress, egress, or network traffic is matched. The action specified in the entry determine how the packets are handled, such as dropping or forwarding.

- Enter a filter entry ID. The system does not dynamically assign a value.

- Assign an action.
- Specify matching criteria.

The following displays an IPv4 filter entry configuration example:

```
A:ALA-7>config>filter>ip-filter# info

description "filter-main"
scope exclusive
entry 10 create
 description "no-91"
 match
 dst-ip 10.10.10.91/24
 src-ip 10.10.0.100/24
 exit
 no action
exit

A:ALA-7>config>filter>ip-filter#
```

### Configuring the HTTP-Redirect Option

If http-redirect is specified as an action, a corresponding forward entry must be specified before the redirect. Http-redirect is not supported on the 7450 ESS-1 model.

The following displays an http-redirect configuration example:

```
A:ALA-48>config>filter>ip-filter# info

description "Captive Portal Filter"
scope template
entry 10 create
 description "Allow DNS"
 match protocol udp
 dst-port eq 53
 exit
 action forward
exit
entry 20 create
 description "Allow Captive Portal"
 match protocol tcp
 dst-ip 100.0.0.2/32
 dst-port eq 80
 exit
 action forward
exit
entry 30 create
 description "HTTP Redirect to Captive Portal"
 match protocol tcp
 dst-port eq 80
 exit
 action http-redirect "http://100.0.0.2/login.cgi?mac=MACsap=$S

```

```

AP&ip=$IP&orig_url=$URL"
 exit

A:ALA-48>config>filter>ip-filter#

```

## Cflowd Filter Sampling

Within a filter entry, you can specify that traffic matching the associated IPv4 filter entry is sampled if the IPv4 interface is set to cflowd acl mode. Enabling **filter-sample** enables the cflowd tool.

The following displays an IPv4 filter entry configuration example:

```

A:ALA-7>config>filter>ip-filter# info

description "filter-main"
scope exclusive
entry 10 create
 description "no-91"
 filter-sample
 interface-disable-sample
 match
 exit
 action forward redirect-policy redirect1
exit

A:ALA-7>config>filter>ip-filter#

```

Within a filter entry, you can also specify that traffic matching the associated IPv4 filter entry is not sampled by cflowd if the IPv4 interface is set to cflowd interface mode. The following displays an IPv4 filter entry configuration example:

```

A:ALA-7>config>filter>ip-filter# info

description "filter-main"
scope exclusive
entry 10 create
 description "no-91"
 no filter-sample
 no interface-disable-sample
 match
 exit
 action forward redirect-policy redirect1
exit

A:ALA-7>config>filter>ip-filter#

```

### 4.2.1.2 Creating an IPv6 Filter Policy

Configuring and applying IPv6 filter policies is optional. IPv6 filter policies must be configured separately from IP (IPv4) filter policies. The configuration mimics IP filter policy configuration. See [Creating an IPv4 Filter Policy](#).

### 4.2.1.3 Creating a MAC Filter Policy

Configuring and applying filter policies is optional. Each filter policy must have the following:

- The filter policy type specified (MAC normal, MAC isid, MAC vid)
- A filter policy ID
- A default action, either drop or forward
- Filter policy scope, either exclusive or template
- At least one filter entry, with a match criterion defined

#### 4.2.1.3.1 MAC Filter Policy

The following example shows a MAC filter policy configuration:

```
A:ALA-7>config>filter# info

...
 mac-filter 90 create
 description "filter-west"
 scope exclusive
 type normal
 exit

A:ALA-7>config>filter#
```

#### 4.2.1.3.2 MAC ISID Filter Policy

The following example shows an ISID filter policy configuration:

```
A;ALA-7>config>filter# info

mac-filter 90 create
 description "filter-wan-man"
 scope template
 type isid
 entry 1 create
```

```

 description "drop-local-isids"
 match
 isid 100 to 1000
 exit
 action drop
 exit
 entry 2 create
 description "allow-wan-isids"
 match
 isid 150
 exit
 action forward
 exit

```

#### 4.2.1.3.3 MAC VID Filter Policy

The following example shows a VID filter policy configuration:

```

A:TOP_NODE>config>filter>mac-filter# info

 default-action forward
 type vic
 entry 1 create
 match frame-type ethernet_II
 ouiter-tag 85 4095
 exit
 action drop
 exit
 entry 2 create
 match frame-type ethernet_II
 ouiter-tag 43 4095
 exit
 action drop
 exit

A:TOP_NODE>config>filter>mac-filter#

```

#### 4.2.1.3.4 MAC Filter Entry

Within a filter policy, configure filter entries that contain criteria against which ingress, egress, or network traffic is matched. The action specified in the entry determines how the packets are handled, such as dropping or forwarding.

- Enter a filter entry ID. The system does not dynamically assign a value.
- Assign an action.
- Specify matching criteria.

The following displays a MAC filter entry configuration example:

```

A:sim1>config>filter# info

```

```

 mac-filter 90 create
 entry 1 create
 description "allow-104"
 match
 exit
 action drop
 exit
 exit

A:sim1>config>filter#

```

#### 4.2.1.4 Creating an IP Exception Filter Policy

Configuring and applying IP exception filter policies is optional. Each exception filter policy must have the following:

- an exception filter policy ID
- scope specified, either exclusive or template
- at least one filter entry with matching criteria specified

##### 4.2.1.4.1 IP Exception Filter Policy

Use the following CLI syntax to create an IP exception filter policy:

**CLI Syntax:**    config>filter# ip-exception *filter-id* [create]  
                               description *description-string*  
                               scope {exclusive | template}

**Example:**        config>filter# ip-exception 1 create  
                       config>filter>ip-except# description "IP-exception"  
                       config>filter>ip-except# scope template

The following example displays a template IP exception filter policy configuration.

```

A:domain1>config>filter# info

...
 ip-exception 1 create
 description "IP-exception"
 scope template
 exit
...

A:domain1>config>filter#

```



#### 4.2.1.4.2 IP Exception Entry Matching Criteria

Within an exception filter policy, configure exception entries that contain criteria against which ingress, egress, and network traffic is matched. Packets that match the entry criteria are allowed to transit the NGE domain in clear text.

- Enter an exception filter entry ID. The system does not dynamically assign a value.
- Specify matching criteria.

Use the following CLI syntax to configure IP exception filter matching criteria:

**CLI Syntax:**

```
config>filter# ip-exception filter-id
 entry entry-id [create]
 description description-string
 match
 dst-ip {ip-address/mask | ip-address ipv4-
 address-mask}
 dst-port {lt | gt | eq} dst-port-number
 dst-port range dst-port-number dst-port-
 number
 icmp-code icmp-code
 icmp-type icmp-type
 src-ip {ip-address/mask | ip-address ipv4-
 address-mask}
 src-port {lt | gt | eq} src-port-number
 src-port range src-port-number src-port-
 number
```

**Example:**

```
config>filter>ip-except# entry 1 create
config>filter>ip-except>entry# match
config>filter>ip-except>entry>match# src-ip 10.10.10.10/
32
config>filter>ip-except>entry>match# dst-ip 10.10.10.91/
24
config>filter>ip-except>entry>match# exit
```

The following example displays a matching configuration.

```
A:domain1>config>filter>ip-exception# info

description "exception-main"
scope exclusive
entry 1
 match
 dst-ip 10.10.10.91/24
 src-ip 10.10.10.10/32
 exit
exit

```

```
A:domain1>config>filter>ip-except#
```

### 4.2.1.5 Creating a Match List for Filter Policies

IP filter policies support usage of match lists as a single match criteria. To create a match list you must:

- Specify a type of a match list (IPv4 address prefix for example).
- Define a unique match list name (IPv4PrefixBlacklist for example).
- Specify at least one list argument (a valid IPv4 address prefix for example).

Optionally, a description can also be defined.

The following example shows an IPv4 address prefix list configuration and its usage in an IPv4 filter policy:

```
*A:ala-48>config>filter# info

match-list
 ip-prefix-list "IPv4PrefixBlacklist"
 description "default IPv4 prefix blacklist"
 prefix 10.0.0.0/21
 prefix 10.254.0.0/24
 exit
exit
ip-filter 10
 scope template
 filter-name "IPv4PrefixBlacklistFilter"
 entry 10
 match
 src-ip ip-prefix-list IPv4PrefixBlacklist
 exit
 action drop
 exit
exit

```

### 4.2.1.6 Applying Filter Policies

Filter policies can be associated with the entities listed in [Table 47](#).

**Table 47** Applying Filter Policies

| IPv4 and IPv6 Filter Policies | MAC Filter Policies  |
|-------------------------------|----------------------|
| Epipe SAP, spoke SDP          | Epipe SAP, spoke SDP |

**Table 47 Applying Filter Policies (Continued)**

| IPv4 and IPv6 Filter Policies                          | MAC Filter Policies           |
|--------------------------------------------------------|-------------------------------|
| Fpipe SAP, spoke SDP                                   | N/A                           |
| IES interface SAP, spoke SDP, R-VPLS                   | N/A                           |
| lpipe SAP, spoke SDP                                   | N/A                           |
| VPLS mesh SDP, spoke SDP, SAP                          | VPLS mesh SDP, spoke SDP, SAP |
| VPRN interface SAP, spoke SDP, R-VPLS, network ingress | N/A                           |
| Network interface                                      | N/A                           |

#### 4.2.1.6.1 Applying IPv4/IPv6 and MAC Filter Policies to a Service

IP and MAC filter policies are applied by associating them with a SAP and/or spoke-sdp in ingress and/or egress direction as needed. Filter ID is used to associate an existing filter policy, or if defined, a Filter Name for that filter policy can be used in the CLI.

The following output displays IP and MAC filters assigned to an ingress and egress SAP and spoke SDP:

```
A:ALA-48>config>service>epipe# info

 sap 1/1/1.1.1 create
 ingress
 filter ip 10
 exit
 egress
 filter mac 92
 exit
 exit
 spoke-sdp 8:8 create
 ingress
 filter ip "epipe sap default filter"
 exit
 egress
 filter mac 91
 exit
exit
no shutdown

A:ALA-48>config>service>epipe#
```

The following output displays an IPv6 filters assigned to an IES service interface:

```
A:ALA-48>config>service>ies# info

 interface "testA" create
```

```

 address 192.22.1.1/24
 sap 2/1/3:0 create
 exit
 ipv6
 ingress
 filter ipv6 100
 egress
 filter ipv6 100
 exit
 exit
...

A:ALA-48>config>service>ies#

```

#### 4.2.1.6.2 Applying IPv4/IPv6 Filter Policies to a Network Port

IP filter policies can be applied to network IPv4/IPv6 interfaces. MAC filters cannot be applied to network IP interfaces or to routable IES services. Similarly to applying filter policies to service, IPv4/IPv6 filter policies are applied to network interfaces by associating a policy with ingress and/or egress direction as desired. Filter ID is used to associate an existing filter policy, or if defined, a Filter Name for that Filter ID policy can be used in the CLI.

The following displays an IP filter applied to an interface at ingress.

```

A:ALA-48>config>router# info
#-----
IP Configuration
#-----
...
 interface "to-104"
 address 10.0.0.103/24
 port 1/1/1
 ingress
 filter ip 10
 exit
 egress
 filter ip "default network egress policy"
 exit
 exit
...
#-----
A:ALA-48>config>router#

```

The following displays IPv4 and IPv6 filters applied to an interface at ingress and egress.

```

A:config>router>if# info

 port 1/1/1
 ipv6
 address 3FFE::101:101/120
 exit

```

```

 ingress
 filter ip 2
 filter ipv6 1
 exit
 egress
 filter ip 2
 filter ipv6 1
 exit

A:config>router>if#

```

### 4.2.1.7 Creating a Redirect Policy

Configuring and applying redirect policies is optional. Each redirect policy must have the following:

- A destination IP address
- A priority (default is 100)
- At least one of the following tests must be enabled:
  - Ping test
  - SNMP test
  - URL test

The following displays a redirection policy configuration:

```

A:ALA-7>config>filter# info

 redirect-policy "redirect1" create
 destination 10.10.10.104 create
 description "SNMP_to_104"
 priority 105
 snmp-test "SNMP-1"
 interval 30
 drop-count 30 hold-down 120
 exit
 no shutdown
 exit
 destination 10.10.10.105 create
 priority 95
 ping-test
 timeout 30
 drop-count 5
 exit
 no shutdown
 exit
 destination 10.10.10.106 create
 priority 90
 url-test "URL_to_106"
 url "http://aww.alcatel.com/ipd/"
 interval 60
 return-code 2323 4567 raise-priority 96

```

---

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

```

A:ALA-7>config>filter#
```

## 4.3 Filter Management Tasks

This section describes filter policy management tasks.

### 4.3.1 Renumbering Filter Policy Entries

The system exits the matching process when the first match is found and then executes the actions in accordance with the specified action. Because the ordering of entries is important, the numbering sequence may need to be rearranged. Entries should be numbered from the most explicit to the least explicit.

The following example shows renumbering of filter entries.

**Example:**

```
config>filter>ip-filter# renum 10 15
config>filter>ip-filter# renum 20 10
config>filter>ip-filter# renum 40 1
```

The following displays the original filter entry order, followed by the reordered filter entries:

```
A:ALA-7>config>filter# info

...
 ip-filter 11 create
 description "filter-main"
 scope exclusive
 entry 10 create
 description "no-91"
 filter-sample
 interface-disable-sample
 match
 dst-ip 10.10.10.91/24
 src-ip 10.10.10.103/24
 exit
 action forward redirect-policy redirect1
 exit
 entry 20 create
 match
 dst-ip 10.10.10.91/24
 src-ip 10.10.0.100/24
 exit
 action drop
 exit
 entry 30 create
 match
 dst-ip 10.10.10.91/24
 src-ip 10.10.0.200/24
 exit
 action forward
```

```
 exit
 entry 40 create
 match
 dst-ip 10.10.10.91/24
 src-ip 10.10.10.106/24
 exit
 action drop
 exit
exit
...

A:ALA-7>config>filter#

A:ALA-7>config>filter# info

...
 ip-filter 11 create
 description "filter-main"
 scope exclusive
 entry 1 create
 match
 dst-ip 10.10.10.91/24
 src-ip 10.10.10.106/24
 exit
 action drop
 exit
 entry 10 create
 match
 dst-ip 10.10.10.91/24
 src-ip 10.10.0.100/24
 exit
 action drop
 exit
 entry 15 create
 description "no-91"
 filter-sample
 interface-disable-sample
 match
 dst-ip 10.10.10.91/24
 src-ip 10.10.10.103/24
 exit
 action forward redirect-policy
 redirect1
 exit
 entry 30 create
 match
 dst-ip 10.10.10.91/24
 src-ip 10.10.0.200/24
 exit
 action forward
 exit
exit
...

A:ALA-7>config>filter#
```



## 4.3.2 Modifying a Filter Policy

There are several ways to modify an existing filter policy. A filter policy can be modified dynamically as part of subscriber management dynamic insertion/removal of filter policy entries (see the *7450 ESS, 7750 SR, and VSR Triple Play Service Delivery Architecture Guide* for details). A filter policy can be modified indirectly by configuration change to a match list the filter policy uses (as described earlier in this guide). In addition, a filter policy can be directly edited as described below.

To access a specific IP (v4/v6), or MAC filter, you must specify the filter ID, or if defined, filter name. Use the **no** form of the command to remove the command parameters or return the parameter to the default setting.

**Example:**

```
config>filter>ip-filter# description "New IP filter
info"
config>filter>ip-filter# entry 2 create
config>filter>ip-filter>entry$ description "new entry"
config>filter>ip-filter>entry# action drop
config>filter>ip-filter>entry# match dst-ip
10.10.10.104/32
config>filter>ip-filter>entry# exit
config>filter>ip-filter#
```

The following output displays the modified IP filter output:

```
A:ALA-7>config>filter# info

...
ip-filter 11 create
description "New IP filter info"
scope exclusive
entry 1 create
match
dst-ip 10.10.10.91/24
src-ip 10.10.10.106/24
exit
action drop
exit
entry 2 create
description "new entry"
match
dst-ip 10.10.10.104/32
exit
action drop
exit
entry 10 create
match
dst-ip 10.10.10.91/24
src-ip 10.10.0.100/24
exit
action drop
```

```

 exit
 entry 15 create
 description "no-91"
 match
 dst-ip 10.10.10.91/24
 src-ip 10.10.10.103/24
 exit
 action forward
 exit
 entry 30 create
 match
 dst-ip 10.10.10.91/24
 src-ip 10.10.0.200/24
 exit
 action forward
 exit
exit
..

A:ALA-7>config>filter#

```

### 4.3.3 Deleting a Filter Policy

Before you can delete a filter, you must remove the filter association from all the applied ingress and egress SAPs and network interfaces by executing **no filter** command in all context where the filter is used.

**Example:**

```

config>service# epipe 5
config>service>epipe# sap 1/1/2:3
config>service>epipe>sap# ingress
config>service>epipe>sap>ingress# no filter

```

After you have removed the filter from the SAPs network interfaces, you can delete the filter as shown in the following example.

**Example:**

```

config>filter# no ip-filter 11

```

### 4.3.4 Modifying a Redirect Policy

To access a specific redirect policy, you must specify the policy name. Use the **no** form of the command to remove the command parameters or return the parameter to the default setting.

**Example:**

```

config>filter# redirect-policy redirect1

```

```
config>filter>redirect-policy# description "New redirect
info"
config>filter>redirect-policy# destination 10.10.10.106
config>filter>redirect-policy>dest# no url-test
"URL_to_106"
config>filter>redirect-policy>dest# url-test
"URL_to_Proxy"
config>filter>redirect-policy>dest>url-test$ url http://
www.alcatel.com
config>filter>redirect-policy>dest>url-test# interval 10
config>filter>redirect-policy>dest>url-test# timeout 10
config>filter>redirect-policy>dest>url-test# return-
code 1
4294967295 raise-priority 255
```

A:ALA-7>config>filter# info

```

...
redirect-policy "redirect1" create
description "New redirect info"
destination 10.10.10.104 create
description "SNMP_to_104"
priority 105
snmp-test "SNMP-1"
interval 30
drop-count 30 hold-down 120
exit
no shutdown
exit
destination 10.10.10.105 create
priority 95
ping-test
timeout 30
drop-count 5
exit
no shutdown
exit
destination 10.10.10.106 create
priority 90
url-test "URL_to_Proxy"
url "http://www.alcatel.com"
interval 10
timeout 10
return-code 1 4294967295 raise-priority 255
exit
no shutdown
exit
no shutdown
exit
...

A:ALA-7>config>filter#
```

### 4.3.5 Deleting a Redirect Policy

Before you can delete a redirect policy from the filter configuration, you must remove the policy association from the IP filter.

The following example shows the command usage to replace the configured redirect policy (**redirect1**) with a different redirect policy (**redirect2**) and then removing the **redirect1** policy from the filter configuration.

**Example:**

```
config>filter>ip-filter 11
config>filter>ip-filter# entry 1
config>filter>ip-filter>entry# action forward redirect-
policy redirect2
config>filter>ip-filter>entry# exit
config>filter>ip-filter# exit
config>filter# no redirect-policy redirect1
```

```
A:ALA-7>config>filter>ip-filter# info

description "This is new"
scope exclusive
entry 1 create
 filter-sample
 interface-disable-sample
 match
 dst-ip 10.10.10.91/24
 src-ip 10.10.10.106/24
 exit
 action forward redirect-policy redirect2
exit
entry 2 create
 description "new entry"
...

A:ALA-7>config>filter>ip-filter#
```

### 4.3.6 Copying Filter Policies

When changes are to be made to an existing filter policy applied to a one or more SAPs/network interfaces, Nokia recommends to first copy the applied filter policy, then modify the copy and then overwrite the applied policy with the modified copy. This ensures that a policy being modified is not applied when partial changes are done as any filter policy edits are applied immediately to all services where the policy is applied.

New filter policies can also be created by copying an existing policy and renaming the new filter.

The following displays the command usage to copy an existing IP filter (**11**) to create a new filter policy (**12**) that can then be edited. And once edits are completed, it can be used to overwrite existing policy (**11**).

**Example:**      config>filter# copy ip-filter 11 to 12

```
A:ALA-7>config>filter# info

...
 ip-filter 11 create
 description "This is new"
 scope exclusive
 entry 1 create
 match
 dst-ip 10.10.10.91/24
 src-ip 10.10.10.106/24
 exit
 action drop
 exit
 entry 2 create
...
 ip-filter 12 create
 description "This is new"
 scope exclusive
 entry 1 create
 match
 dst-ip 10.10.10.91/24
 src-ip 10.10.10.106/24
 exit
 action drop
 exit
 entry 2 create
...

A:ALA-7>config>filter#
```



## 4.4 Filter Configuration Command Reference

- [Command Hierarchies](#)
- [Command Descriptions](#)

### 4.4.1 Command Hierarchies

- [IPv4 Filter Policy Commands](#)
- [IPv6 Filter Policy Commands](#)
- [MAC Filter Commands](#)
- [IP Exception Filter Policy Configuration Commands](#)
- [System Filter Policy Commands](#)
- [Redirect Policy Configuration Commands](#)
- [Match Filter List Commands](#)
- [Log Filter Commands](#)
- [Copy Filter Commands](#)

#### 4.4.1.1 IPv4 Filter Policy Commands

```

config
 — filter
 — ip-filter filter-id [create]
 — no ip-filter {filter-id | filter-name}
 — [no] chain-to-system-filter
 — default-action {drop | forward}
 — description description-string
 — no description
 — embed-filter filter-id [offset offset] [{active | inactive}]
 — no embed-filter filter-id
 — embed-filter flowspec [group group-id] [router {router-instance | service-

 name vprn-service-name}] [offset offset] [{active | inactive}]
 — no embed-filter flowspec [group group-id]
 — embed-filter open-flow ofs-name [{system | service {service-id | service-

 name} | sap sap-id}] [offset offset] [{active | inactive}]
 — no embed-filter open-flow ofs-name [{system | service {service-id | service-

 name} | sap sap-id}]
 — embed-filter vsd vsd-filter-id [offset offset] [{active | inactive}]
 — no embed-filter vsd vsd-filter-id
 — entry entry-id [create]
 — no entry entry-id
 — [no] action [secondary]

```

- **drop**
- **drop** **packet-length** {lt | gt | eq} *packet-length-value*
- **drop** **packet-length range** *packet-length-value packet-length-value*
- **drop** **ttl** {lt | gt | eq} *ttl-value*
- **drop range** *ttl-value ttl-value*
- [no] **extended-action**
  - **remark dscp** *dscp-name*
- **forward**
- **forward** **bonding-connection** *connection-id*
- **forward** **esi esi sf-ip** *ip-address vas-interface interface-name router router-instance*
- **forward** **esi esi sf-ip** *ip-address vas-interface interface-name router service-name service-name*
- **forward** **esi esi service-id** *vpls-service-id*
- **forward** **lsp** *lsp-name*
- **forward** **next-hop** *ip-address*
- **forward** **next-hop** *ip-address router router-instance*
- **forward** **next-hop** *ip-address router service-name service-name*
- **forward** **next-hop indirect** *ip-address*
- **forward** **next-hop indirect** *ip-address router router-instance*
- **forward** **next-hop indirect** *ip-address router service-name service-name*
- **forward** **next-hop interface** *ip-int-name*
- **forward** **redirect-policy** *policy-name*
- **forward** **router** *router-instance*
- **forward** **router** *service-name service-name*
- **forward** **sap** *sap-id*
- **forward** **sdp** *sdp-id:vc-id*
- **forward** **vprn-target bgp-nh** *ip-address router router-instance [adv-prefix ip-address/prefix-length] [lsp lsp-name]*
- **forward** **vprn-target bgp-nh** *ip-address router service-name service-name [adv-prefix ip-address/prefix-length] [lsp lsp-name]*
- **gtp-local-breakout**
- **http-redirect** *rdr-url-string [allow-radius-override]*
- **nat** [**nat-policy** *nat-policy-name*]
- **rate-limit** *value*
- **rate-limit** *value* **packet-length** {{lt | eq | gt} *packet-length-value*
- **rate-limit** *value* **packet-length range** *packet-length-value packet-length-value*
- **rate-limit** *value* **ttl** {lt | gt | eq} *ttl-value*
- **rate-limit** *value* **ttl range** *ttl-value ttl-value*
- **reassemble**
- **remark**
- **tcp-mss-adjust**
- **description** *description-string*
- **no description**
- **egress-pbr** {**default-load-balancing** | **l4-load-balancing**}
- **no egress-pbr**
- [no] **filter-sample**
- [no] **interface-disable-sample**
- **log** *log-id*



- **no log**
- **match** [*protocol protocol-id*]
- **no match**
  - **dscp** *dscp-name*
  - **no dscp**
  - **dst-ip** {*ip-address/mask* | *ip-address ipv4-address-mask* | *ip-prefix-list prefix-list-name*}
  - **no dst-ip**
  - **dst-port** {*lt* | *gt* | *eq*} *dst-port-number*
  - **dst-port** *port-list* *port-list-name*
  - **dst-port range** *dst-port-number dst-port-number*
  - **no dst-port**
  - **fragment** {*true* | *false*}
  - **no fragment**
  - **icmp-code** *icmp-code*
  - **no icmp-code**
  - **icmp-type** *icmp-type*
  - **no icmp-type**
  - **ip-option** *ip-option-value* [*ip-option-mask*]
  - **no ip-option**
  - **multiple-option** {*true* | *false*}
  - **no multiple-option**
  - **option-present** {*true* | *false*}
  - **no option-present**
  - **port** {*lt* | *gt* | *eq*} *port-number*
  - **port** *port-list* *port-list-name*
  - **port range** *port-number port-number*
  - **no port**
  - **src-ip** {*ip-address/mask* | *ip-address ipv4-address-mask* | *ip-prefix-list prefix-list-name*}
  - **no src-ip**
  - **src-port** {*lt* | *gt* | *eq*} *src-port-number*
  - **src-port** *port-list* *port-list-name*
  - **src-port range** *src-port-number src-port-number*
  - **no src-port**
  - **src-route-option** {*true* | *false*}
  - **no src-route-option**
  - **tcp-ack** {*true* | *false*}
  - **no tcp-ack**
  - **tcp-syn** {*true* | *false*}
  - **no tcp-syn**
- **pbr-down-action-override** {*drop* | *forward* | *filter-default-action*}
- **no pbr-down-action-override**
- **sticky-dest** *hold-time-up*
- **sticky-dest no-hold-time-up**
- **no sticky-dest**
- **filter-name** *filter-name*
- **no filter-name**
- **renum** *old-entry-id new-entry-id*
- **scope** {*exclusive* | *template* | *embedded* | *system*}
- **no scope**
- **shared-radius-filter-wmark** *low low-watermark high high-watermark*
- **no shared-radius-filter-wmark**
- **sub-insert-credit-control** *start-entry entry-id count count*

- **no sub-insert-credit-control**
- **sub-insert-radius** *start-entry entry-id count count*
- **no sub-insert-radius**
- **sub-insert-shared-pccrule** *start-entry entry-id count count*
- **no sub-insert-shared-pccrule**
- **sub-insert-shared-radius** *start-entry entry-id count count*
- **no sub-insert-shared-radius**
- **sub-insert-wmark** *low low-watermark high high-watermark*
- **no sub-insert-wmark**

#### 4.4.1.2 IPv6 Filter Policy Commands

These commands do not apply to the 7450 ESS (except in mixed mode).

```

config
— filter
— ipv6-filter filter-id [create]
— no ipv6-filter {filter-id | filter-name}
— [no] chain-to-system-filter
— default-action {drop | forward}
— description description-string
— no description
— embed-filter filter-id [offset offset] [{active | inactive}]
— no embed-filter filter-id
— embed-filter flowspec [group group-id] [router {router-instance | service-
name vprn-service-name}] [offset offset] [{active | inactive}]
— no embed-filter flowspec [group group-id]
— embed-filter open-flow ofs-name [{system | service {service-id | service-
name} | sap sap-id}] [offset offset] [{active | inactive}]
— no embed-filter open-flow ofs-name [{system | service {service-id | service-
name} | sap sap-id}]
— embed-filter vsd vsd-filter-id [offset value] [{active | inactive}]
— no embed-filter vsd vsd-filter-id
— entry entry-id [create]
— no entry entry-id
— [no] action [secondary]
— drop
— drop hop-limit {lt | gt | eq} hop-limit-value
— drop hop-limit range hop-limit-value hop-limit-value
— drop payload-length {lt | gt | eq} payload-length-value
— drop payload-length range payload-length-value payload-
length-value
— [no] extended-action
— remark dscp dscp-name
— forward
— forward bonding-connection connection-id
— forward esi esi sf-ip ip-address vas-interface interface-name
router router-instance
— forward esi esi sf-ip ip-address vas-interface interface-name
router service-name service-name
— forward esi esi service-id vpls-service-id

```

- **forward** *lsp* *lsp-name*
- **forward** **next-hop** *ip-address*
- **forward** **next-hop** *ip-address* **router** *router-instance*
- **forward** **next-hop** *ip-address* **router** **service-name** *service-name*
- **forward** **next-hop indirect** *ip-address*
- **forward** **next-hop indirect** *ip-address* **router** *router-instance*
- **forward** **next-hop indirect** *ip-address* **router** **service-name** *service-name*
- **forward** **redirect-policy** *policy-name*
- **forward** **router** *router-instance*
- **forward** **router** **service-name** *service-name*
- **forward** **sap** *sap-id*
- **forward** **sdp** *sdp-id:vc-id*
- **forward** **vprn-target** **bgp-nh** *ip-address* **router** *router-instance* [**adv-prefix** *ip-address/prefix-length*] [**lsp** *lsp-name*]
- **forward** **vprn-target** **bgp-nh** *ip-address* **router** **service-name** *service-name* [**adv-prefix** *ip-address/prefix-length*] [**lsp** *lsp-name*]
- **http-redirect** *rdr-url-string* [**allow-radius-override**]
- **nat** [**nat-policy** *nat-policy-name*]
- **rate-limit** *value*
- **rate-limit** *value* **hop-limit** {**lt** | **eq** | **gt**} *hop-limit-value*
- **rate-limit** *value* **hop-limit range** *hop-limit-value* *hop-limit-value*
- **rate-limit** *value* **payload-length** {**lt** | **gt** | **eq**} *payload-length-value*
- **rate-limit** *value* **payload-length range** *payload-length-value* *payload-length-value*
- **remark** **dscp** *dscp-name*
- **tcp-mss-adjust**
- **description** *description-string*
- **no description**
- **egress-pbr** {**default-load-balancing** | **l4-load-balancing**}
- **no egress-pbr**
- [**no**] **filter-sample**
- [**no**] **interface-disable-sample**
- **log** *log-id*
- **no log**
- **match** [**next-header** *next-header*]
- **no match**
  - **ah-ext-hdr** {**true** | **false**}
  - **no ah-ext-hdr**
  - **dscp** *dscp-name*
  - **no dscp**
  - **dst-ip** {*ipv6-address/prefix-length* | *ipv6-address* *ipv6-address-mask* | **ipv6-prefix-list** *prefix-list-name*}
  - **no dst-ip**
  - **dst-port** {**lt** | **gt** | **eq**} *dst-port-number*
  - **dst-port** **port-list** *port-list-name*
  - **dst-port range** *dst-port-number* *dst-port-number*
  - **no dst-port**
  - **esp-ext-hdr** {**true** | **false**}
  - **no esp-ext-hdr**
  - **flow-label** *flow-label* [*mask*]

- 
- **no flow-label**
  - **fragment** {true | false | first-only | non-first-only}
  - **no fragment**
  - **hop-by-hop-opt** {true | false}
  - **no hop-by-hop-opt**
  - **icmp-code** *icmp-code*
  - **no icmp-code**
  - **icmp-type** *icmp-type*
  - **no icmp-type**
  - **port** {lt | gt | eq} *port-number*
  - **port port-list** *port-list-name*
  - **port range** *port-number port-number*
  - **no port**
  - **routing-type0** {true | false}
  - **no routing-type0**
  - **src-ip** {*ipv6-address/prefix-length* | *ipv6-address ipv6-address-mask* | **ipv6-prefix-list** *prefix-list-name*}
  - **no src-ip**
  - **src-port** {lt | gt | eq} *src-port-number*
  - **src-port port-list** *port-list-name*
  - **src-port range** *src-port-number src-port-number*
  - **no src-port**
  - **tcp-ack** {true | false}
  - **no tcp-ack**
  - **tcp-syn** {true | false}
  - **no tcp-syn**
  - **pbr-down-action-override** {drop | forward | filter-default-action}
  - **no pbr-down-action-override**
  - **sticky-dest** *hold-time-up*
  - **sticky-dest no-hold-time-up**
  - **no sticky-dest**
  - **filter-name** *filter-name*
  - **no filter-name**
  - **renum** *old-entry-id new-entry-id*
  - **scope** {exclusive | template | embedded | system}
  - **no scope**
  - **shared-radius-filter-wmark** **low** *low-watermark* **high** *high-watermark*
  - **no shared-radius-filter-wmark**
  - **sub-insert-credit-control** **start-entry** *entry-id* **count** *count*
  - **no sub-insert-credit-control**
  - **sub-insert-radius** **start-entry** *entry-id* **count** *count*
  - **no sub-insert-radius**
  - **sub-insert-shared-pccrule** **start-entry** *entry-id* **count** *count*
  - **no sub-insert-shared-pccrule**
  - **sub-insert-shared-radius** **start-entry** *entry-id* **count** *count*
  - **no sub-insert-shared-radius**
  - **sub-insert-wmark** **low** *low-watermark* **high** *high-watermark*
  - **no sub-insert-wmark**

### 4.4.1.3 MAC Filter Commands

```

config
 — filter
 — mac-filter filter-id [create]
 — mac-filter {filter-id | filter-name}
 — no mac-filter {filter-id | filter-name}
 — default-action {drop | forward}
 — description description-string
 — no description
 — embed-filter vsd vsd-filter-id [offset value] [{active | inactive}]
 — no embed-filter vsd vsd-filter-id
 — entry entry-id [create]
 — no entry entry-id
 — [no] action [secondary]
 — drop
 — forward
 — forward esi esi service-id vppls-service-id
 — forward sap sap-id
 — forward sdp sdp-id:vc-id
 — rate-limit value
 — description description-string
 — no description
 — log log-id
 — no log
 — match [frame-type {802dot3 | 802dot2-llc | 802dot2-snap |
 ethernet_II}]
 — no match
 — dot1p dot1p-value [dot1p-mask]
 — no dot1p
 — dsap dsap-value [dsap-mask]
 — no dsap
 — dst-mac ieee-address [ieee-address-mask]
 — no dst-mac
 — etype 0x0600..0xffff
 — no etype
 — inner-tag value [vid-mask]
 — no inner-tag
 — isid value [to higher-value]
 — no isid
 — outer-tag value [vid-mask]
 — no outer-tag
 — snap-oui {zero | non-zero}
 — no snap-oui
 — snap-pid snap-pid
 — no snap-pid
 — ssap ssap-value [ssap-mask]
 — no ssap
 — src-mac ieee-address [ieee-address-mask]
 — no src-mac
 — pbr-down-action-override {drop | forward | filter-default-action}
 — no pbr-down-action-override
 — sticky-dest hold-time-up

```

- **sticky-dest** no-hold-time-up
- no **sticky-dest**
- **filter-name** *filter-name*
- no **filter-name**
- **renum** *old-entry-id new-entry-id*
- **scope** {exclusive | template}
- no **scope**
- **type** *filter-type*

#### 4.4.1.4 IP Exception Filter Policy Configuration Commands

- ```

config
  — filter
    — ip-exception filter-id [create]
    — [no] ip-exception {filter-id | filter-name}
      — description description-string
      — no description
      — entry entry-id [create]
      — no entry entry-id
        — description description-string
        — no description
        — match [protocol protocol-id]
        — no match
          — dst-ip {ip-address/mask | ip-address ipv4-address-mask}
          — no dst-ip
          — dst-port {lt | gt | eq} dst-port-number
          — dst-port range dst-port-number dst-port-number
          — no dst-port
          — icmp-code icmp-code
          — no icmp-code
          — icmp-type icmp-type
          — no icmp-type
          — src-ip {ip-address/mask | ip-address ipv4-address-mask}
          — no src-ip
          — src-port {lt | gt | eq} src-port-number
          — src-port range src-port-number src-port-number
          — no src-port
      — filter-name filter-name
      — no filter-name
      — renum old-entry-id new-entry-id
      — scope {exclusive | template}
      — no scope

```

4.4.1.5 System Filter Policy Commands

- ```

config
 — filter
 — system-filter

```

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

#### 4.4.1.6 Redirect Policy Configuration Commands

```

config
 — filter
 — redirect-policy redirect-policy-name [create]
 — no redirect-policy redirect-policy-name
 — description description-string
 — no description
 — destination ip-address [create]
 — no destination ip-address
 — description description-string
 — no description
 — [no] ping-test
 — drop-count consecutive-failures [hold-down seconds]
 — no drop-count
 — interval seconds
 — no interval
 — timeout seconds
 — no timeout
 — priority [priority]
 — no priority
 — [no] shutdown
 — snmp-test test-name [create]
 — no snmp-test test-name
 — drop-count consecutive-failures [hold-down seconds]
 — no drop-count
 — interval seconds
 — no interval
 — oid oid-string community community-string
 — no oid
 — return-value return-value type return-type [disable | lower-
 priority priority | raise-priority priority]
 — no return-value return-value type return-type
 — timeout seconds
 — no timeout
 — [no] unicast-rt-test
 — url-test test-name [create]
 — no url-test test-name
 — drop-count consecutive-failures [hold-down seconds]
 — no drop-count
 — interval seconds
 — no interval
 — return-code return-code-1 [return-code-2] [disable | lower-
 priority priority | raise-priority priority]
 — no return-code return-code-1 [return-code-2]
 — timeout seconds
 — no timeout
 — url url-string [http-version version-string]

```

- **no url**
- **router** *router-instance*
- **router service-name** *service-name*
- **no router**
- **[no] shutdown**
- **sticky-dest** *hold-time-up*
- **sticky-dest no-hold-time-up**
- **no sticky-dest**

#### 4.4.1.7 Match Filter List Commands

- ```

config
  — filter
    — match-list
      — ip-prefix-list ip-prefix-list-name [create]
      — no ip-prefix-list ip-prefix-list-name
        — [no] apply-path
          — bgp-peers criterion-index group reg-exp neighbor reg-exp
          — bgp-peers criterion-index router router-instance group reg-exp
            neighbor reg-exp
          — bgp-peers criterion-index router service-name service-name
            group reg-exp neighbor reg-exp
          — no bgp-peers criterion-index
        — description description-string
        — no description
        — [no] prefix ip-prefix/prefix-length
      — ipv6-prefix-list ipv6-prefix-list-name [create]
      — no ipv6-prefix-list ipv6-prefix-list-name
        — [no] apply-path
          — bgp-peers criterion-index group reg-exp neighbor reg-exp
          — bgp-peers criterion-index router router-instance group reg-exp
            neighbor reg-exp
          — bgp-peers criterion-index router service-name service-name
            group reg-exp neighbor reg-exp
          — no bgp-peers criterion-index
        — description description-string
        — no description
        — [no] prefix ipv6-prefix/prefix-length
      — port-list port-list-name [create]
      — no port-list port-list-name
        — description description-string
        — no description
        — [no] port port-number
        — [no] port range start end

```

4.4.1.8 Log Filter Commands

```

config

```



```

— filter
  — log log-id [create]
  — no log log-id
    — description description-string
    — no description
    — destination {memory num-entries | syslog syslog-id}
    — no destination
    — [no] shutdown
    — summary
      — [no] shutdown
      — summary-crit dst-addr
      — summary-crit src-addr
      — no summary-crit
    — [no] wrap-around

```

4.4.1.9 Copy Filter Commands

```

config
  — filter
    — copy
      — ip-filter src-filter-id [src-entry src-entry-id] to dst-filter-id [dst-entry dst-entry-id] [overwrite]
      — mac-filter src-filter-id [src-entry src-entry-id] to dst-filter-id [dst-entry dst-entry-id] [overwrite]
      — ipv6-filter src-filter-id [src-entry src-entry-id] to dst-filter-id [dst-entry dst-entry-id] [overwrite]

```

4.4.2 Command Descriptions

- [Generic Commands](#)
- [Global Filter Commands](#)
- [Filter Log Commands](#)
- [ACL Filter Policy Commands](#)
- [General Filter Entry Commands](#)
- [IP \(v4/v6\) and IP Exception Filter Entry Commands](#)
- [Match List Configuration Commands](#)
- [MAC Filter Entry Commands](#)
- [MAC Filter Match Criteria](#)
- [Policy and Entry Maintenance Commands](#)
- [Redirect Policy Commands](#)

4.4.2.1 Generic Commands

description

| | |
|--------------------|---|
| Syntax | description <i>description-string</i> no description |
| Context | config>filter>ip-exception config>filter>ip-exception>entry config>filter>ip-filter config>filter>ipv6-filter config>filter>ip-filter>entry config>filter>ipv6-filter>entry config>filter>log config>filter>mac-filter config>filter>mac-filter>entry config>filter>redirect-policy config>filter>redirect-policy>destination config>filter>match-list>ip-prefix-list config>filter>match-list>ipv6-prefix-list config>filter>match-list>port-list |
| Description | <p>This command creates a text description stored in the configuration file for a configuration context.</p> <p>The no form of the command removes any description string from the context.</p> |
| Default | no description |
| Parameters | <i>description-string</i> — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. |

4.4.2.2 Global Filter Commands

ip-filter

| | |
|--------------------|--|
| Syntax | ip-filter <i>filter-id</i> [create] no ip-filter { <i>filter-id</i> <i>filter-name</i> } |
| Context | config>filter |
| Description | This command creates a configuration context for the specified IPv4 filter policy. |

The **no** form of the command deletes the IPv4 filter policy. A filter policy cannot be deleted until it is removed from all objects where it is applied.

- Default** No IPv4 filter policy is created by default.
- Parameters** *filter-id* — Specifies the IPv4 filter policy ID expressed as a decimal integer.
- Values** 1 to 65535
- create** — This keyword is required to create the configuration context. Once it is created, the context can be enabled with or without the **create** keyword.
- filter-name* — Specifies a string of up to 64 characters uniquely identifying this IPv4 filter policy.

ipv6-filter

- Syntax** **ipv6-filter** *filter-id* [**create**]
no ipv6-filter {*filter-id* | *filter-name*}
- Context** config>filter
- Description** This command creates a configuration context for the specified IPv6 filter policy.
- The **no** form of the command deletes the IPv6 filter policy. A filter policy cannot be deleted until it is removed from all objects where it is applied.
- Default** No IPv6 filter policy is created by default.
- Parameters** *filter-id* — Specifies the IPv6 filter policy ID expressed as a decimal integer.
- Values** 1 to 65535
- create** — This keyword is required to create the configuration context. Once it is created, the context can be enabled with or without the **create** keyword.
- filter-name* — Specifies a string of up to 64 characters uniquely identifying this IPv6 filter policy.

system-filter

- Syntax** **system-filter**
- Context** config>filter
- Description** This command enables the context to activate system filter policies.

mac-filter

| | |
|--------------------|---|
| Syntax | mac-filter <i>filter-id</i> [create] mac-filter { <i>filter-id</i> <i>filter-name</i> } no mac-filter { <i>filter-id</i> <i>filter-name</i> } |
| Context | config>filter |
| Description | <p>This command, creates a configuration context for the specified MAC filter policy.</p> <p>The no form of the command deletes the MAC filter policy. A filter policy cannot be deleted until it is removed from all objects where it is applied.</p> |
| Parameters | <p><i>filter-id</i> — Specifies the MAC filter policy ID expressed as a decimal integer.</p> <p>Values 1 to 65535</p> <p>create — keyword required to create the configuration context. Once it is created, the context can be enabled with or without the create keyword.</p> <p><i>filter-name</i> — a string of up to 64 characters uniquely identifying this MAC filter policy</p> |

redirect-policy

| | |
|--------------------|--|
| Syntax | redirect-policy <i>redirect-policy-name</i> [create] no redirect-policy <i>redirect-policy-name</i> |
| Context | config>filter |
| Description | <p>This command, creates a configuration context for the specified redirect policy.</p> <p>The no form of the command removes the redirect policy from the filter configuration only if the policy is not referenced in a filter and the filter is not in use (applied to a service or network interface).</p> |
| Parameters | <p><i>redirect-policy-name</i> — Specifies the redirect policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. There is no limit to the number of redirect policies that can be configured.</p> <p>create — This keyword is required to create the configuration context. Once it is created, the context can be enabled with or without the create keyword.</p> |

log

| | |
|---------------|---|
| Syntax | log <i>log-id</i> [create] no log <i>log-id</i> |
|---------------|---|

| | |
|----------------------|---|
| Context | config>filter |
| Description | <p>This command, creates a configuration context for the specified filter log if it does not exist, and enables the context to configure the specified filter log.</p> <p>The no form of the command deletes the filter log. The log cannot be deleted if there are filter entries configured to write to the log. All filter entry logging associations need to be removed before the log can be deleted.</p> |
| Default | log 101 |
| Special Cases | Filter log 101 — Filter log 101 is the default log and is automatically created by the system. Filter log 101 is always a memory filter log and cannot be changed to a Syslog filter log. The log size defaults to 1000 entries, the filter log description is set to "Default filter log". The number of entries and wrap-around behavior can be modified. |
| Parameters | <p><i>log-id</i> — Specifies the filter log ID expressed as a decimal integer</p> <p>Values 101 to 199</p> <p>create — keyword required to create the configuration context. Once it is created, the context can be enabled with or without the create keyword.</p> |

4.4.2.3 Filter Log Commands

destination

| | |
|--------------------|--|
| Syntax | destination memory <i>num-entries</i> destination syslog <i>syslog-id</i> no destination |
| Context | config>filter>log |
| Description | <p>This command configures the destination for filter log entries for the filter log ID.</p> <p>Filter logs can be sent to either memory (memory) or to an existing Syslog server definition (syslog).</p> <p>If the filter log destination is memory, the maximum number of entries in the log must be specified.</p> <p>The no form of the command deletes the filter log association.</p> |
| Default | destination memory 1000 |

| | |
|-------------------|---|
| Parameters | memory <i>num-entries</i> — Specifies the destination of the filter log ID is a memory log. The <i>num-entries</i> value is the maximum number of entries in the filter log expressed as a decimal integer. Values 10 to 50000 |
| | syslog <i>syslog-id</i> — Specifies the destination of the filter log ID is a Syslog server. The <i>syslog-id</i> parameter is the number of the Syslog server definition. Values 1 to 10 |

shutdown

| | |
|--------------------|--|
| Syntax | [no] shutdown |
| Context | config>filter>log config>filter>log>summary |
| Description | <p>Administratively enables/disables (AdminUp/AdminDown) an entity. Downing an entity does not change, reset or remove any configuration settings or statistics. Many objects must be shutdown before they may be deleted.</p> <p>The shutdown command administratively downs an entity. Administratively downing an entity changes the operational state of the entity to down.</p> <p>Unlike other commands and parameters where the default state will not be indicated in the configuration file, shutdown and no shutdown are always indicated in system generated configuration files.</p> <p>The no form of the command puts an entity into the administratively enabled state.</p> |
| Default | no shutdown (for config>filter>log) shutdown (for config>filter>log>summary) |

summary

| | |
|--------------------|---|
| Syntax | summary |
| Context | config>filter>log |
| Description | This command enables the context to configure log summarization. These settings will only be taken into account when syslog is the log destination. |
| Default | n/a |

summary-crit

| | |
|--------------------|---|
| Syntax | summary-crit dst-addr summary-crit src-addr no summary-crit |
| Context | config>filter>log>summary |
| Description | <p>This command defines the key of the index of the minitable. If key information is changed while summary is administratively enabled (no shutdown), the filter summary minitable is flushed and recreated with different key information. Log packets received during the reconfiguration time will be handled as if summary was not active.</p> <p>The no form of the command reverts to the default parameter.</p> |
| Default | summary-crit src-addr |
| Parameters | <p>dst-addr — Specifies that received log packets are summarized based on the destination IPv4, IPv6, or MAC address.</p> <p>src-addr — Specifies that received log packets are summarized based on the source IPv4, IPv6 or MAC address</p> |

wrap-around

| | |
|--------------------|--|
| Syntax | [no] wrap-around |
| Context | config>filter>log |
| Description | <p>This command configures a memory filter log to log until full or to store the most recent log entries (circular buffer).</p> <p>Specifying wrap-around configures the memory filter log to store the most recent filter log entries (circular buffer). When the log is full, the oldest filter log entries are overwritten with new entries.</p> <p>The no form of the command configures the memory filter log to accept filter log entries until full. When the memory filter log is full, filter logging for the log filter ID ceases.</p> |
| Default | wrap-around |

4.4.2.4 ACL Filter Policy Commands

default-action

| | |
|---------------|--|
| Syntax | default-action {drop forward} |
|---------------|--|

| | |
|--------------------|---|
| Context | config>filter>ip-filter config>filter>ipv6-filter config>filter>mac-filter |
| Description | This command defines the default action to be applied to packets not matching any entry in this ACL filter policy or to packets for that match a PBF/PBR filter entry for which the PBF/PBR target is down and pbr-down-action-override per-entry is set to filter-default-action . |
| Default | default-action drop |
| Parameters | drop — Specifies the default action is to drop a packet. forward — Specifies the default action is to forward a packet. |

chain-to-system-filter

| | |
|--------------------|---|
| Syntax | [no] chain-to-system-filter |
| Context | config>filter>ip-filter config>filter>ipv6-filter |
| Description | <p>This command chains this filter to a currently active system filter. When the filter is chained to the system filter, the system filter rules are executed first, and the filter rules are only evaluated if no match on the system filter was found.</p> <p>The no form of the command detaches this filter from the system filter.</p> <p>Operational note:</p> <p>If no system filter is currently active, the command has no effect.</p> |
| Default | no chain-to-system-filter |

ip

| | |
|--------------------|---|
| Syntax | [no] ip <i>ip-filter-id</i> |
| Context | config>filter>system-filter |
| Description | <p>This command activates an IPv4 system filter policy. Once activated, all IPv4 ACL filter policies that chain to the system filter (config>filter>ip-filter>chain-to-system-filter) will automatically execute system filter policy rules first.</p> <p>The no form of the command deactivates the system filter policy.</p> |
| Parameters | <p><i>ip-filter-id</i> — Specifies the an existing IPv4 filter policy with scope system. This parameter can either be expressed as a decimal integer, or as an ASCII string of up to 64 characters.</p> <p>Values 1 to 65535 or the filter policy name (<i>filter-name</i>, 64 char max)</p> |

ipv6

| | |
|--------------------|--|
| Syntax | [no] ipv6 <i>ipv6-filter-id</i> |
| Context | config>filter>system-filter |
| Description | This command activates an IPv6 system filter policy. Once activated, all IPv6 ACL filter policies that chain to the system filter (config>filter>ipv6-filter>chain-to-system-filter) will automatically execute system filter policy rules first. The no form of the command deactivates the system filter policy. |
| Parameters | <i>ipv6-filter-id</i> — Specifies the an existing IPv6 filter policy with scope system . This parameter can either be expressed as a decimal integer, or as an ASCII string of up to 64 characters in length. Values 1 to 65535 or the filter policy name |

embed-filter

| | |
|--------------------|---|
| Syntax | embed-filter <i>filter-id</i> [offset <i>offset</i>] [{ active inactive }] no embed-filter <i>filter-id</i> embed-filter flowspec [group <i>group-id</i>] [router { <i>router-instance</i> service-name <i>vprn-service-name</i> }] [offset <i>value</i>] [{ active inactive }] no embed-filter flowspec [group <i>group-id</i>] embed-filter open-flow <i>ofs-name</i> [{ system service { <i>service-id</i> <i>service-name</i> }] sap <i>sap-id</i>] [offset <i>offset</i>] [{ active inactive }] no embed-filter open-flow <i>ofs-name</i> [{ system service { <i>service-id</i> <i>service-name</i> }] sap <i>sap-id</i>] embed-filter vsd <i>vsd-filter-id</i> [offset <i>value</i>] [{ active inactive }] no embed-filter vsd <i>vsd-filter-id</i> |
| Context | config>filter>ip-filter config>filter>ipv6-filter config>filter>mac-filter |
| Description | This command embeds a previously defined IPv4, IPv6, or MAC embedded filter policy or Hybrid OpenFlow switch instance into this exclusive, template, or system filter policy at the specified offset value. Rules derived from the BGP flowspec can also be embedded into template filter policies only. |



Note: For MAC filters, embedding is supported for VSD filters or filter entries only.

The **embed-filter open-flow** *ofs-name* form of this command enables OpenFlow (OF) in GRT either by embedding the specified OpenFlow switch (OFS) instance with **switch-defined-cookie** disabled, or by embedding rules with `sros-cookie:type "grt-cookie"`, value 0, from the specified OFS instance with **switch-defined-cookie** enabled. The embedding filter can only be deployed in GRT context or be unassigned.

The **embed-filter open-flow** *ofs-name system* form of this command enables OF in system filters by embedding rules with `sros-cookie:type "system-cookie"`, value 0, from the specified OFS instance with **switch-defined-cookie** enabled. The embedding filter can only be of scope **system**.

The **embed-filter open-flow** *ofs-name service* {*service-id* | *service-name*} form of this command enables OF in VPRN/VPLS filters by embedding rules with `sros-cookie:type "service-cookie"`, value **service-id**, from the specified OFS instance with **switch-defined-cookie** enabled—per service rules. The embedding filter can only be deployed in the specified VPRN/VPLS service. A single VPLS service can only support OF rules per SAP or per service.

The **embed-filter open-flow** *ofs-name sap* *sap-id* form of this command enables OF in VPLS SAP filters by embedding rules with `sros-cookie:type "service-cookie"`, value *service-id* and flow match conditions specifying the *sap-id* from the specified OFS instance with **switch-defined-cookie** enabled—per SAP OF rules. The embedding filter must be of type **exclusive** and can only be deployed on the specified SAP in the context of the specified VPLS service. A single VPLS service can only support OF rules per SAP or per service.

The **no embed-filter open-flow** *ofs-name* form of this command removes the OF embedding for the GRT context.

The **embed-filter flowspec** form of this command enables the embedding of rules derived from BGP flowspec routes into the filter policy that is being configured. The optional **group** parameter specifies that only flowspec routes tagged with an interface-set extended community containing this group ID should be selected for embedding. The optional **router** parameter specifies the routing instance source of the BGP flowspec routes; if the parameter is not specified, the routing instance is derived automatically from the context in which the filter policy is applied. Flowspec rules associated with one routing instance cannot be embedded in a filter applied to an interface of a different routing instance. After flowspec rules associated with one routing instance are embedded into a filter, that filter policy cannot be applied to an interface of a different routing instance.

The **no embed-filter flowspec** form of this command removes the flowspec filter embedding from this filter policy.

The **embed-filter vsd** *vsd-filter-id* command refers to the VSD filter ID encoded `_tmnx_vsd_filter-id`. The filter is created dynamically and managed exclusively using the Python script, so rules can be inserted and removed in the correct VSD filters. The command is supported with IP, IPv6, and MAC filters. For more information about VSD filter provisioning, automation, and the Python script, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN*.

The **no embed-filter vsd** *vsd-filter-id* form of this command removes the VSD filter embedding from this filter policy.

The **no embed-filter** *filter-id* form of this command removes the embedding from this filter policy.

See the description of embedded filter policies in this guide for further operational details.

Parameters

filter-id — Specifies a previously defined embedded filter policy.

offset — Specifies an embedded filter entry X will have an entry X + offset in the embedding filter.

Values 0 to 65536

Default 0

active — Specifies that embedded filter entries are to be included in this embedding filter policy and activated on applicable line cards—default if no keyword is specified and omitted from **info** command output (but not **info detail**), or when saving the configuration.

inactive — Specifies that no embedded filter policy entries are to be included in this embedding filter policy. The embedding is configured but will not do anything.

flowspec — This keyword indicates that rules derived from BGP flowspec routes should be embedded into (or removed from, in case of the **no** form) the filter.

group-id — Specifies that only flowspec routes with an interface-set extended community with this value of *group-id* should be selected for embedding.

Values 0 to 16383

router-instance — Specifies a router instance.

system — Used for OF control of system filters. Allows embedding of OF rules into system filters from OFS with **switch-defined-cookie** enabled. Only the rules with cookie value encoding “system-cookie” are embedded.

sap-id — Used for OF control of VPLS services when a PortID and VLAN ID match is required. Allows embedding of OF rules with a PortID and VLAN ID match into exclusive VPLS SAP filters. Only the rules with cookie value encoding the VPLS service, and flow table match encoding the specified SAP, are embedded into the filter. The embedding filter can only be deployed in the context of the specified SAP.

sap-id — Specifies an existing SAP that the embedding filter can be used for.

ofs-name — Specifies the name of the currently configured Hybrid OpenFlow Switch (OFS) instance.

Not including the **system**, **service** or **sap** parameters will specify OF in a GRT instance context by default. This allows embedding of OF rules into filters deployed in GRT instances from OFS with **switch-defined-cookie** disabled, or embedding rules from OFS with **switch-defined-cookie** enabled, when the FlowTable cookie encodes sros-cookie:type “grt-cookie”.

{*service-id* | *service-name*} — Used for OF control of VPRN or VPLS services. Allows embedding of OF rules into a VPRN or VPLS access or network filters. Only the rules with cookie value encoding the specified service ID are embedded into the filter. The embedding filter can only be deployed in the context of the specified service.

service-id — Specifies an existing VPRN or VPLS service ID that the embedding filter can be used for.

service-name — Specifies an existing VPRN or VPLS service name that the embedding filter can be used for.

vsd-filter-id — Creates an embedded filter (filter ID: `_tmnx_vsd_filter-id`) for population by Nuage VSD.

filter-name

| | |
|--------------------|--|
| Syntax | filter-name <i>filter-name</i> no filter-name |
| Context | config>filter>ip-exception config>filter>ip-filter config>filter>ipv6-filter config>filter>mac-filter |
| Description | This command configures filter-name attribute of a given filter. filter-name, when configured, can be used instead of filter ID to reference the given policy in the CLI. |
| Default | no filter-name |
| Parameters | <i>filter-name</i> — Specifies a string up to 64 characters in length that uniquely identifies this filter policy. The following restrictions apply to the <i>filter-name</i> : <ul style="list-style-type: none"> • Policy names may not begin with a number (0-9). • Policy names may not begin with the underscore “_” character (e.g. <code>_myPolicy</code>). Names that start with underscore are reserved for system generated names. • “fSpec-x” (where x is any number) cannot be used as a user defined filter name. |

scope

| | |
|--------------------|---|
| Syntax | scope { exclusive template embedded system } scope { exclusive template } no scope |
| Context | config>filter>ip-exception config>filter>ip-filter config>filter>ipv6-filter config>filter>mac-filter |
| Description | This command configures the filter policy scope as exclusive, template, embedded or system. The scope of the policy cannot be changed when: |

- the scope is **template** and the policy is applied to one or more services or network interfaces
- the scope is **embedded** and the policy is embedded by another policy

Changing the scope to/from system is only allowed when a policy is not active and the policy has no entries configured.

The **no** form of the command sets the scope of the policy to the default of **template**.

| | |
|-------------------|--|
| Default | scope template |
| Parameters | <p>exclusive — Specifies that the policy can only be applied to a single entity. Attempting to assign the policy to a second entity will result in an error message.</p> <p>template — Specifies that the policy can be applied to multiple entities.</p> <p>embedded — Specifies that the policy cannot be applied directly. The policy defines embedded filter rules, which are embedded by other exclusive/template/system filter policies. The embedded scope is supported for IPv4 and IPv6 filter policies only.</p> <p>system — Specifies that the policy defines system-wide filter rules. To apply system policy rules, activate system filter and chain exclusive/template ACL filter policy to the system filter. The system scope is supported for IPv4 and IPv6 filter policies only.</p> |

shared-radius-filter-wmark

| | |
|--------------------|---|
| Syntax | shared-radius-filter-wmark low low-watermark high high-watermark no shared-radius-filter-wmark |
| Context | config>filter>ip-filter config>filter>ipv6-filter |
| Description | This command configures the low and high watermark for the number of RADIUS shared filters reporting |
| Default | no shared-radius-filter-wmark |
| Parameters | <p><i>low-watermark</i> — Specifies the utilization of the filter ranges for filter entry insertion, at which a table full alarm will be raised by the agent.</p> <p>Values 0 to 8000</p> <p><i>high-watermark</i> — Specifies the utilization of the filter ranges for filter entry insertion, at which a table full alarm will be cleared by the agent.</p> <p>Values 1 to 8000</p> |

sub-insert-credit-control

| | |
|---------------|---|
| Syntax | sub-insert-credit-control start-entry entry-id count count |
|---------------|---|

no sub-insert-credit-control

| | |
|--------------------|---|
| Context | config>filter>ip-filter config>filter>ipv6-filter |
| Description | This command inserts point information for credit control for the filter. The no form of the command reverts to the default. |
| Default | no sub-insert-credit-control |
| Parameters | <i>entry-id</i> — Identifies a filter on this system. Values 1 to 65535 <i>count count</i> — Specifies the count Values 1 to 65535 |

sub-insert-radius

| | |
|--------------------|--|
| Syntax | sub-insert-radius start-entry <i>entry-id</i> count <i>count</i> no sub-insert-radius |
| Context | config>filter>ip-filter config>filter>ipv6-filter |
| Description | This command inserts point information for RADIUS for the filter. The no form of the command reverts to the default. |
| Default | no sub-insert-radius |
| Parameters | <i>entry-id</i> — Specifies at what place the filter entries received from RADIUS will be inserted in the filter. Values 1 to 65535 <i>count</i> — Specifies the count. Values 1 to 65535 |

sub-insert-shared-pccrule

| | |
|----------------|---|
| Syntax | sub-insert-shared-pccrule start-entry <i>entry-id</i> count <i>count</i> no sub-insert-shared-pccrule |
| Context | config>filter>ip-filter config>filter>ipv6-filter |

| | |
|--------------------|---|
| Description | This command defines the range of filter and QoS policy entries that are reserved for shared entries received in Flow-Information AVP via Gx interface (PCC rules – Policy and Charging Control). The no form of this command disables the insertion, which will result in a failure of PCC rule installation. |
| Default | no sub-insert-shared-pccrule |
| Parameters | <i>entry-id</i> — Specifies the lowest entry in the range. <div style="margin-left: 40px;">Values 1 to 65535</div> <i>count</i> — Specifies the number of entries in the range. <div style="margin-left: 40px;">Values 1 to 65535</div> |

sub-insert-shared-radius

| | |
|--------------------|---|
| Syntax | sub-insert-shared-radius start-entry <i>entry-id</i> count <i>count</i> no sub-insert-shared-radius |
| Context | config>filter>ip-filter config>filter>ipv6-filter |
| Description | This command configures the insert point for shared host rules from RADIUS. |
| Default | no sub-insert-shared-radius |
| Parameters | <i>entry-id</i> — Identifies a filter on this system. <div style="margin-left: 40px;">Values 1 to 65535</div> <i>count</i> — Specifies the count. <div style="margin-left: 40px;">Values 1 to 65535</div> |

sub-insert-wmark

| | |
|--------------------|---|
| Syntax | sub-insert-wmark low <i>low-watermark</i> high <i>high-watermark</i> no sub-insert-wmark |
| Context | config>filter>ip-filter config>filter>ipv6-filter |
| Description | This command configures the low and high watermark percentage for inserted filter entry usage reporting. The no form of the command reverts to the default. |
| Default | sub-insert-wmark low 90 high 95 |

| | |
|-------------------|--|
| Parameters | <i>low-watermark</i> — Specifies the utilization of the filter ranges for filter entry insertion, at which a table full alarm will be cleared by the agent. Values 0 to 100 <i>high-watermark</i> — Specifies the utilization of the filter ranges for filter entry insertion, at which a table full alarm will be raised by the agent. Values 0 to 100 |
|-------------------|--|

type

| | |
|--------------------|---|
| Syntax | type <i>filter-type</i> |
| Context | config>filter>mac-filter |
| Description | This command configures the MAC Filter Policy sub-type as being either normal, ISID or VID. |
| Default | type normal |
| Parameters | <i>filter-type</i> — Specifies which type of entry this MAC filter can contain. Values normal — regular match criteria are allowed; ISID or VID filter match criteria not allowed isid — only ISID match criteria are allowed vid — only VID match criteria are allowed on ethernet_II frame types |

4.4.2.5 General Filter Entry Commands

entry

| | |
|--------------------|--|
| Syntax | entry <i>entry-id</i> [create] no entry <i>entry-id</i> |
| Context | config>filter>ip-exception config>filter>ip-filter config>filter>ipv6-filter config>filter>mac-filter |
| Description | <p>This command creates or edits an IPv4, IPv6, MAC, or IP exception filter entry. Multiple entries can be created using unique <i>entry-id</i> numbers within the filter. Entries must be sequenced from most to least explicit.</p> <p>An entry may not have any match criteria defined (in which case everything matches) but must have at least the keyword action for it to be considered complete. Entries without the action keyword will be considered incomplete and hence will be rendered inactive.</p> |

The **no** form of the command removes the specified entry from the filter. Entries removed from the filter are immediately removed from all services or network ports where that filter is applied.

Parameters *entry-id* — Uniquely identifies a match criteria and the corresponding action. It is recommended that multiple entries be given *entry-id* in staggered increments. This allows users to insert a new entry in an existing policy without requiring to renumber all the existing entries. The parameter is expressed as a decimal integer.

Values 1 to 65535

create — This keyword is required to create the configuration context. Once the context is created, the user can enable the context with or without the **create** keyword.

action

Syntax [no] action [secondary]

Context config>filter>ip-filter>entry
config>filter>ipv6-filter>entry
config>filter>mac-filter>entry

Description This command enters the context to configure a primary (no option specified) or secondary (**secondary** option specified) action to be performed on packets matching this filter entry. An ACL filter entry remains inactive (is not programmed in hardware) until a specific action is configured for that entry.

A primary action supports any filter entry action, a secondary action is used for redundancy and defines a redundant L3 PBR action for an L3 PBR primary action or a redundant L2 PBF action for a L2 PBF primary action.

The **no** form of this command removes the specific action configured in the context of the action command. The primary action cannot be removed if a secondary action exists.

Default no action

Parameters **secondary** — Specifies a secondary action to be performed on packets matching this filter entry. A secondary action can only be configured if a primary action is configured.

log

Syntax log log-id
no log

Context config>filter>ip-filter>entry
config>filter>ipv6-filter>entry
config>filter>mac-filter>entry

| | | | |
|--------------------|--|---------------|------------|
| Description | <p>This command associates a filter log to the current filter policy entry and therefore enables logging for that filter entry.</p> <p>The filter log must exist before a filter entry can be enabled to use the filter log.</p> <p>The no form of the command disables logging for the filter entry.</p> | | |
| Default | no log | | |
| Parameters | <p><i>log-id</i> — Specifies the filter log ID expressed as a decimal integer.</p> <table><tr><td>Values</td><td>101 to 199</td></tr></table> | Values | 101 to 199 |
| Values | 101 to 199 | | |

pbr-down-action-override

| | |
|--------------------|--|
| Syntax | pbr-down-action-override {drop forward filter-default-action} no pbr-down-action-override |
| Context | config>filter>ip-filter>entry config>filter>ipv6-filter>entry config>filter>mac-filter>entry |
| Description | <p>This command allows overriding the default action that is applied for entries with PBR/PBF action defined, when the PBR/PBF target is down.</p> <p>The no form of the command preserves default behavior when PBR/PBF target is down.</p> |
| Default | no pbr-down-action-override |
| Parameters | <p>drop — Specifies that packets matching the entry will be dropped if PBR/PBF target is down.</p> <p>forward — Specifies that packets matching the entry will be forwarded if PBR/PBF target is down.</p> <p>filter-default-action — Specifies that packets matching the entry will be processed as per default-action configuration for this filter if PBR/PBF target is down.</p> |

sticky-dest

| | |
|----------------|--|
| Syntax | sticky-dest hold-time-up sticky-dest no-hold-time-up no sticky-dest |
| Context | config>filter>ip-filter>entry config>filter>ipv6-filter>entry config>filter>mac-filter>entry |

| | |
|--------------------|--|
| Description | <p>This command configures sticky destination behavior for redundant PBR/PBF actions. Configuring sticky destination has an effect on PBR/PBF actions whether a secondary action is configured.</p> <p>The <i>hold-time-up</i> parameter allows the operator to delay programming of a PBR/PBF action for a specified amount of time. The timer is only started when transitioning from all configured targets being down (that is, the primary target if no secondary target is configured, or both the primary and secondary targets when both are configured) to at least one target being up.</p> <p>When the timer expires, the primary PBR/PBF action is programmed if its target is up. If the primary PBR/PBF target is down and a secondary PBR/PBF action has been configured and its target is up, then this secondary PBR/PBF action is programmed. In all other cases, no specific programming occurs when the timer expires.</p> <p>When sticky destination is configured and the secondary PBR/PBF target is up and its associated action is programmed, it is not automatically replaced by the primary PBR/PBF action when its target transitions from down to up. In this situation, programming the primary PBR/PBF action can be forced using the activate-primary-action tools command.</p> <p>Changing the value of the timer while the timer is running takes effect immediately (that is, the timer is restarted immediately using the new value).</p> <p>The no form of the command disables sticky destination behavior.</p> |
| Default | no sticky-dest |
| Parameters | <p><i>hold-time-up</i> — Specifies the initial delay in seconds. Zero is equivalent to no-hold-time-up (no delay).</p> <p>Values 0 to 65535 seconds</p> |

4.4.2.6 IP (v4/v6) and IP Exception Filter Entry Commands

action

| | |
|--------------------|---|
| Syntax | action |
| Context | config>filter>ip-filter>entry config>filter>ip-filter>entry>action |
| Description | This command (under the config>filter>ip-filter>entry context) sets the context for specific action commands to be performed (under the config>filter>ip-filter>entry>action context) on packets matching this filter entry. |

drop

| | |
|--------------------|---|
| Syntax | drop drop packet-length {lt gt eq} <i>packet-length-value</i> drop packet-length range <i>packet-length-value</i> <i>packet-length-value</i> drop ttl {lt gt eq} <i>tll-value</i> drop ttl range <i>tll-value</i> <i>tll-value</i> |
| Context | config>filter>ip-filter>entry>action |
| Description | This command sets the filter entry action to drop. |

drop

| | |
|--------------------|--|
| Syntax | drop drop hop-limit {lt gt eq} <i>hop-limit-value</i> drop hop-limit range hop-limit-value <i>hop-limit-value</i> drop payload-length {lt gt eq} <i>payload-length-value</i> drop payload-length range <i>payload-length-value</i> payload-length-value |
| Context | config>filter>ipv6-filter>entry>action |
| Description | This command sets the filter entry action to drop. |

drop-extracted-traffic

| | |
|--------------------|--|
| Syntax | drop-extracted-traffic |
| Context | config>filter>ip-filter>entry>action config>filter>ipv6-filter>entry>action |
| Description | This command specifies that a packet matching this filter entry is dropped if extracted to the CPM. Packets matching the filter entry match criteria and not extracted to the CPM are forwarded with no further match in following filter entries. |

extended-action

| | |
|----------------|--|
| Syntax | [no] extended-action |
| Context | config>filter>ip-filter>entry>action config>filter>ipv6-filter>entry>action |

Description This command enables the context to configure an extended action for a filter entry's PBR action (configured under **config>filter>ip-filter>entry>action** and **config>filter>ipv6-filter>entry>action** context). The extended action is executed in addition to the configured PBR action.

The **no** form of the command removes the extended action.

Default No extended action is configured by default.

remark

Syntax **remark dscp** *dscp-name*

Context
config>filter>ip-filter>entry>action
config>filter>ip-filter>entry>action>extended-action
config>filter>ipv6-filter>entry>action>extended-action

Description This command enables and configures the remarking of the DiffServ Code Points of packets matching the criteria of the IPv4/IPv6 filter policy entry, in conjunction with a PBR action. Packets are remarked regardless of QoS-based in-profile or out-of-profile classification. QoS-based DSCP remarking is overridden. If the status of the PBR target is tracked and it is down, the extended action will not be executed; otherwise, the extended action will be performed.

Parameters *dscp-name* — Specifies the DSCP value to write

Values be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63

forward

Syntax **forward**
forward bonding-connection *connection-id*
forward esi esi sf-ip *ip-address* **vas-interface** *interface-name* **router** *router-instance*
forward esi esi sf-ip *ip-address* **vas-interface** *interface-name* **router** **service-name** *service-name*
forward esi esi service-id *vpls-service-id*
forward lsp *lsp-name*
forward next-hop *ip-address*
forward next-hop *ip-address* **router** *router-instance*
forward next-hop *ip-address* **router** **service-name** *service-name*
forward next-hop indirect *ip-address*
forward next-hop indirect *ip-address* **router** *router-instance*

forward next-hop indirect *ip-address* **router** **service-name** *service-name*
forward next-hop interface *ip-int-name*
forward redirect-policy *policy-name*
forward router *router-instance*
forward router **service-name** *service-name*
forward sap *sap-id*
forward sdp *sdp-id:vc-id*
forward vprn-target bgp-nh *ip-address* **router** *router-instance* [**adv-prefix** *ip-address/prefix-length*] [**lsp** *lsp-name*]
forward vprn-target bgp-nh *ip-address* **router** **service-name** *service-name* [**adv-prefix** *ip-address/prefix-length*] [**lsp** *lsp-name*]

| | |
|--------------------|---|
| Context | config>filter>ip-filter>entry>action |
| Description | This command sets the context for specific forward commands to be performed. |
| Parameters | <p><i>connection-id</i> — Specifies that the packet should be forwarded over the specified connection (specified by the connection ID under the bonding group interface), if that connect is available. Outside of a bonding egress context, the behavior of this filter is undefined.</p> <p>Values 1, 2</p> <p>esi service-id — Specifies that the packet matching the entry is forwarded to ESI identified first appliance in Nuage service chain using EVPN-resolved VXLAN tunnel in the specified VPLS service.</p> <p>esi sf-ip vas-interface router — Specifies that the packet matching the entry is forwarded to ESI/SF-IP identified first appliance in Nuage service chain using EVPN-resolved VXLAN tunnel over the configured VAS interface in the specified VPRN service.</p> <p>lsp — Specifies that the packet matching the entry is forwarded using the specified lsp.</p> <p>next-hop — Specifies that the packet matching the entry is forwarded in the routing context of the incoming interface using direct or indirect IPv4 address in the routing lookup.</p> <p>next-hop router — Specifies that the packet matching the entry is forwarded in the configured routing context using direct or indirect IPv4 address in the routing lookup.</p> <p>next-hop interface — Specifies that the packet matching the entry is forwarded using the configured local interface.</p> <p>redirect-policy — Specifies that the packet matching the entry is forwarded using forward next-hop or forward nexthop router and the IP address of destination selected by the configured redirect policy.</p> <p>If no destination is selected, packets are subject to action forward.</p> <p>router — Specifies that the packet matching the entry is routed in the configured routing instance and not in the incoming interface routing instance.</p> |

sap — Specifies that the packet matching the entry is forwarded using the configured SAP.

sdp — Specifies that the packet matching the entry is forwarded using the configured SDP.

vprn-target — Specifies that the packet matching the entry is redirected towards a designated BGP next-hop (**bgp-nh**). The user may specify an LSP (**lsp** *lsp-name*) to use towards that next-hop. If no LSP is specified, the system will automatically select one. The user must specify the routing context (**router** {*router-instance* | **service-name** *service-name*}) in which the system will perform the lookups in order to derive the proper VPRN service label. The user may specify an advertised prefix route (**adv-prefix** *ip-address/prefix-length*). This is needed in case label per VRF is not the label allocation method configured at the BGP peer.

esi — Specifies a 10-byte Ethernet Segment Identifier

ip-address/prefix-length — Specifies an IPv4 advertised route in the CIDR notation. The IPv4 address is in dotted decimal notation

Values ip-address d.d.d.d, where “d” is [0...255]D
 prefix-length: 1 to 32

bgp-nh *ip-address* — Specifies the IPv4 address (in dotted decimal notation) of the target BGP next-hop.

Values ip-address d.d.d.d

ip-address — Specifies the IPv4 address of a direct or indirect next hop to which to forward matching packets.

ip-int-name — Specifies the name of an egress IP interface where matching packets will be forwarded from. This parameter is only valid for unnumbered point-to-point interfaces. If the string contains special characters (such as #, \$, spaces), the entire string must be enclosed within double quotes.

interface-name — Specifies the (maximum 32-character) name of an egress R-VPLS IP interface used to forward the packets using ESI redirect for VPRN/IES service.

lsp-name — Specifies an existing RSVP-TE, MPLS-TP, or SR-TE LSP that supports LSP redirect.

policy-name — Specifies an IPv4 redirect policy configured in the config>filter>redirect-policy context.

sap-id — Specifies an existing VPLS Ethernet SAP.

sdp-id:vc-id — Specifies an existing VPLS SDP.

router-instance — Specifies “Base” or an existing VPRN service ID For the **forward vprn-target bgp-nh** command, *router-instance* must specify an existing VPRN service ID.

service-name — Specifies an existing VPRN service name.

vpls-service-id — Specifies an existing VPLS service ID or service name.

forward

| | |
|--------------------|---|
| Syntax | forward forward bonding-connection <i>connection-id</i> forward esi esi sf-ip <i>ipv6-address</i> vas-interface <i>interface-name</i> router <i>router-instance</i> forward esi esi sf-ip <i>ipv6-address</i> vas-interface <i>interface-name</i> router service-name <i>service-name</i> forward esi esi service-id <i>vpls-service-id</i> forward lsp <i>lsp-name</i> forward next-hop <i>ipv6-address</i> forward next-hop <i>ipv6-address</i> router <i>router-instance</i> forward next-hop <i>ipv6-address</i> router service-name <i>service-name</i> forward next-hop indirect <i>ipv6-address</i> forward next-hop indirect <i>ipv6-address</i> router <i>router-instance</i> forward next-hop indirect <i>ipv6-address</i> router service-name <i>service-name</i> forward redirect-policy <i>policy-name</i> forward router <i>router-instance</i> forward router service-name <i>service-name</i> forward sap <i>sap-id</i> forward sdp <i>sdp-id:vc-id</i> forward vprn-target bgp-nh <i>ipv6-address</i> router <i>router-instance</i> [adv-prefix <i>ipv6-address/prefix-length</i>] [lsp <i>lsp-name</i>] forward vprn-target bgp-nh <i>ipv6-address</i> router service-name <i>service-name</i> [adv-prefix <i>ipv6-address/prefix-length</i>] [lsp <i>lsp-name</i>] |
| Context | config>filter>ipv6-filter>entry>action |
| Description | This command sets the context for specific forward commands to be performed. |
| Parameters | <p><i>connection-id</i> — Specifies that the packet be forwarded over the specified connection if that connection is available. The connection is specified by the connection ID under the bonding group interface. Outside of a bonding egress context, the behavior of this filter is undefined.</p> <p>Values 1, 2</p> <p>esi service-id — Specifies that a packet matching the entry is forwarded to an ESI-identified first appliance in the Nuage service chain using an EVPN-resolved VXLAN tunnel in the specified VPLS service.</p> <p>esi sf-ip vas-interface router — Specifies that a packet matching the entry is forwarded to an ESI/SF-IP-identified first appliance in the Nuage service chain using an EVPN-resolved VXLAN tunnel over the configured VAS interface in the specified VPRN service.</p> <p>lsp — Specifies that a packet matching the entry is forwarded using the specified LSP.</p> |

next-hop — Specifies that a packet matching the entry is forwarded in the routing context of the incoming interface using a direct or indirect IPv6 address in the routing lookup.

next-hop router — Specifies that a packet matching the entry is forwarded in the configured routing context using a direct or indirect IPv6 address in the routing lookup.

next-hop interface — Specifies that a packet matching the entry is forwarded using the configured local interface.

redirect-policy — Specifies that a packet matching the entry is forwarded using a forward next-hop or forward nexthop router and the IP address of the destination selected by the configured redirect policy.

If no destination is selected, packets are subject to action **forward**.

router — Specifies that a packet matching the entry is routed in the configured routing instance and not in the incoming interface routing instance.

sap — Specifies that a packet matching the entry is forwarded using the configured SAP.

sdp — Specifies that a packet matching the entry is forwarded using the configured SDP.

vpn-target — Specifies that a packet matching the entry is redirected towards a designated BGP next-hop (**bgp-nh**). The user can specify an LSP (**lsp lsp-name**) to use towards that next-hop. If no LSP is specified, the system automatically select one. The user must specify the routing context (**router {router-instance | service-name service-name}**) in which the system performs the lookups to derive the proper VPRN service label. The user can specify an advertised prefix route (**adv-prefix ip-address/prefix-length**). This is needed if the label per VRF is not the label allocation method configured at the BGP peer.

esi — Specifies a 10-byte Ethernet Segment Identifier.

ipv6-address/prefix-length — Specifies an IPv6 advertised route in the CIDR notation.

bgp-nh ipv6-address — Specifies the IPv6 address of the target BGP next-hop.

ipv6-address — Specifies the IPv6 address of a direct or indirect next hop to forward matching packets in hex digits.

Values

| | |
|----|-------------------------------------|
| x: | x:x:x:x:x:x:x (eight 16-bit pieces) |
| x: | x:x:x:x:x:d.d.d.d |
| x: | [0..FFFF]H |
| d: | [0..255]D |

interface-name — Specifies the interface name (up to 32 characters in length) of an egress R-VPLS IP interface used to forward the packets using ESI redirect for a VPRN or IES service.

lsp-name — Specifies an existing RSVP-TE, MPLS-TP, or SR-TE LSP that supports LSP redirect.

policy-name — Specifies an IPv6 redirect policy configured in the config>filter>redirect-policy context.

sap-id — Specifies an existing VPLS Ethernet SAP.

sdp-id:vc-id — Specifies an existing VPLS SDP.

router-instance — Specifies the router name or CPM router instance.

Values

router-instance : *router name*

router-name Base | management | cpm-vr-name

cpm-vr-name [32 characters maximum]

service-name — Specifies an existing VPRN service name.

vpls-service-id — Specifies an existing VPLS service ID or service name.

gtp-local-breakout

| | |
|--------------------|---|
| Syntax | gtp-local-breakout |
| Context | config>filter>ip-filter>entry |
| Description | This command specifies the filter entry action to gtp-local-breakout. |

http-redirect

| | |
|--------------------|--|
| Syntax | http-redirect |
| Context | config>filter>ip-filter>entry config>filter>ipv6-filter>entry |
| Description | This command sets the filter entry action to http-redirect. |

nat

| | |
|--------------------|--|
| Syntax | nat [nat-policy <i>nat-policy-name</i>] |
| Context | config>filter>ip-filter>entry>action config>filter>ipv6-filter>entry>action |
| Description | <p>This command enables NAT traffic diversion based on IPv4 filters (LSN44) or IPv6 filters (DS-Lite, NAT64). The filter contains a matching condition based on any combination of the 5 tuple. Traffic will be diverted to NAT based on such defined matching condition. Filter fields outside of the 5 tuples are not valid and it will be ignored in filter based traffic diversion to NAT.</p> <p>The pool selection for the outside IP address and port along with other mapping characteristics can be specified by the means on the NAT policy.</p> |

Parameters *nat-policy-name* — Specifies the NAT policy name up to 32 characters in length.

rate-limit

Syntax **rate-limit** *value*
rate-limit *value* **packet-length** {**lt** | **gt** | **eq**} *packet-length-value*
rate-limit *value* **packet-length range** *packet-length-value* *packet-length-value*
rate-limit *value* **ttl** {**lt** | **gt** | **eq**} *ttl-value*
rate-limit *value* **ttl range** *ttl-value* *ttl-value*

Context config>filter>ip-filter>entry>action

Description This command sets the rate limit for the traffic matching both the filter entry match criteria and the *packet-length-value* defined in the **rate-limit action** statement.

Packets matching the filter entry match criteria and not matching the *packet-length-value* defined in the **rate-limit action** statement are implicitly forwarded with no further match in subsequent filter entries.

Rate limit packets matching both the filter entry match criteria and the *ttl-value* are defined in the **action rate-limit** statement.

Packets matching the filter entry match criteria and not matching the *ttl-value* defined in the **rate-limit action** statement are implicitly forwarded with no further match in following filter entries.

Parameters *value* — Specifies the **rate-limit value** in kb/s. A rate of 0 results in all traffic being dropped. A rate of **max** results in all traffic being forwarded.

Values 0 to 2000000000 kb/s, max

packet-length-value — Specifies the packet length value for the rate limit action.

Values 0 to 65523

packet-length — Specifies that rate-limit packets matching both the filter entry match criteria and the *packet-length value* defined in the **rate-limit** action statement. Packets matching the filter entry match criteria and not matching the *packet-length* value defined in the **rate-limit** action statement are implicitly forwarded with no further match in following filter entries.

Values **lt** — Specifies “less than”. The **lt** parameter cannot be used with the lowest possible numerical value for the parameter.

eq — Specifies “equal to”.

gt — Specifies “greater than”. The **gt** parameter cannot be used with the highest possible numerical value for the parameter.

range — Specifies an inclusive range. When **range** is used, the start of the range (the first value entered) must be smaller than the end of the range (the second value entered).

ttl-value — Specifies rate-limit packets matching both the filter entry match criteria and the TTL value defined in the *rate-limit* action statement. Packets matching the filter entry match criteria and not matching the TTL value defined in the *rate-limit* action statement are implicitly forwarded with no further match in following filter entries.

Values 0 to 255

rate-limit

Syntax **rate-limit** *value*
rate-limit *value* **hop-limit** {*lt* | *gt* | *eq*} *hop-limit-value*
rate-limit *value* **hop-limit range** *hop-limit-value* *hop-limit-value*
rate-limit *value* **payload-length** {*lt* | *gt* | *eq*} *payload-length-value*
rate-limit *value* **payload-length range** *payload-length-value* *payload-length-value*

Context config>filter>ipv6-filter>entry>action

Description This command sets rate-limit packets matching both the filter entry match criteria and the *payload-length-value* defined in the **rate-limit action** statement.

Packets matching the filter entry match criteria and not matching the *payload-length-value* defined in the **rate-limit action** statement are implicitly forwarded with no further match in following filter entries.

Rate limit packets matching both the filter entry match criteria and the *hop-limit-value* are defined in the **rate-limit action** statement.

Packets matching the filter entry match criteria and not matching the *hop-limit-value* defined in the **action rate-limit** statement are implicitly forwarded with no further match in following filter entries.

Parameters *value* — Specifies the **rate-limit** *value* in kb/s. A rate of 0 results in all traffic being dropped. A rate of **max** results in all traffic being forwarded.

Values 0 to 2000000000 kb/s | max

hop-limit — Specifies the hop limit value for the rate limit action.

Values *lt* — Specifies “less than”. The *lt* parameter cannot be used with the lowest possible numerical value for the parameter.
eq — Specifies “equal to”.
gt — Specifies “greater than”. The *gt* parameter cannot be used with the highest possible numerical value for the parameter.

hop-limit-value — Specifies the hop limit value for the rate limit action.

Values 0 to 255

payload-length-value — Specifies the payload length value for the rate limit action.

Values 0 to 65535

payload-length — Specifies rate-limit packets matching both the filter entry match criteria and the *payload-length-value* defined in the **rate-limit** action statement. Packets matching the filter entry match criteria and not matching the *payload-length-value* defined in the **rate-limit** action statement are implicitly forwarded with no further match in following filter entries.

Values

- lt — Specifies “less than”. The **lt** parameter cannot be used with the lowest possible numerical value for the parameter.
- eq — Specifies “equal to”.
- gt — Specifies “greater than”. The **gt** parameter cannot be used with the highest possible numerical value for the parameter.

range — Specifies an inclusive range. When the **range** parameter is used, the start of the range (the first value entered) must be smaller than the end of the range (the second value entered).

reassemble

| | |
|--------------------|--|
| Syntax | reassemble |
| Context | config>filter>ip-filter>entry>action config>filter>ipv6-filter>entry>action |
| Description | This command sets the filter entry action to reassemble. |

remark

| | |
|--------------------|--|
| Syntax | remark |
| Context | config>filter>ip-filter>entry>action config>filter>ipv6-filter>entry>action |
| Description | This command activates DSCP remarking for packets matching the entry. |

tcp-mss-adjust

| | |
|--------------------|---|
| Syntax | tcp-mss-adjust |
| Context | config>filter>ip-filter>entry>action config>filter>ipv6-filter>entry>action |
| Description | This command activates adjustment of maximum segment size (MSS) option of TCP packets matching the entry. |

egress-pbr

| | |
|--------------------|---|
| Syntax | egress-pbr {default-load-balancing I4-load-balancing} no egress-pbr |
| Context | config>filter>ip-filter>entry config>filter>ipv6-filter>entry |
| Description | <p>This command specifies that the configured PBR action is applicable to egress processing. The command should only be enabled in ACL policies used by residential subscribers. Enabling egress-pbr on filters not deployed for residential subscribers is not blocked but may lead to unexpected behavior and should be avoided.</p> <p>The no form of this command removes the egress-pbr designation of the filter entry's action.</p> |
| Default | no egress-pbr |
| Parameters | default-load-balancing — Sets load-balancing to the default (hash based on SA/DA of the packet) I4-load-balancing — Includes TCP/UDP port (if available) in the hash |

filter-sample

| | |
|--------------------|---|
| Syntax | [no] filter-sample |
| Context | config>filter>ip-filter>entry config>filter>ipv6-filter>entry |
| Description | <p>This command enables cflowd sampling for packets matching this filter entry.</p> <p>If the cflowd is either not enabled or set to cflowd interface mode, this command is ignored.</p> <p>The no form disables the cflowd sampling using this filter entry.</p> |
| Default | no filter-sample |

interface-disable-sample

| | |
|--------------------|---|
| Syntax | [no] interface-disable-sample |
| Context | config>filter>ip-filter>entry config>filter>ipv6-filter>entry |
| Description | <p>This command disables cflowd sampling for packets matching this filter entry, for the IP interface set to cflowd interface mode. This allows the option to not sample specific types of traffic when interface sampling is enabled.</p> <p>If the cflowd is either not enabled or set to cflowd acl mode, this command is ignored.</p> |

The **no** form of this command enables sampling.

Default no interface-disable-sample

match

Syntax **match** [**protocol** *protocol-id*]
no match

Context config>filter>ip-filter>entry
config>filter>ip-exception>entry

Description This command enables the context to enter match criteria for the filter entry. When the match criteria have been satisfied the action associated with the match criteria is executed.

A **match** context may consist of multiple match criteria, but multiple **match** statements cannot be entered per entry. More precisely, the command can be entered multiple times but this only results in modifying the *protocol-id*, and does not affect the underlying match criteria configuration.

The **no** form of the command removes all the match criteria from the filter entry and sets the *protocol-id* of the match command to **none** (keyword). As per above, **match protocol none** is however not equivalent to **no match**.

Default match next-header none

Parameters *protocol-id* — Configures an IP protocol to be used as an IP filter match criterion. The protocol type, such as TCP or UDP, is identified by its respective protocol number.

protocol-number — Specifies the protocol number.

Values [0..255]D
[0x0..0xFF]H
[0b0..0b11111111]B

protocol-name — Configures the decimal value representing the IP protocol to be used as an IP filter match criterion. Well-known protocol numbers include ICMP(1), TCP(6), UDP(17).

Default Value: none (keyword)

Values 0 to 255 in decimal format. Values can also be specified in hexadecimal format, in binary format, or using the following keywords:

IPv4 filter keywords: none, icmp, igmp, ip, tcp, egp, igp, udp, rdp, ipv6, ipv6-route, ipv6-frag, idrp, rsvp, gre, ipv6-icmp, ipv6-no-nxt, ipv6-opts, iso-ip, eigrp, ospf-igp, ether-ip, encap, pnni, pim, vrrp, l2tp, stp, ptp, isis, crtp, crudp, sctp

IP exception filter keywords: none, icmp, igmp, ospf-igp, pim, rsvp, tcp, udp, vrrp

* — udp/tcp wildcard

Table 48 Protocol ID Descriptions

| Protocol | Protocol ID | Description |
|-------------|-------------|---|
| icmp | 1 | Internet Control Message |
| igmp | 2 | Internet Group Management |
| ip | 4 | IP in IP (encapsulation) |
| tcp | 6 | Transmission Control |
| egp | 8 | Exterior Gateway Protocol |
| igp | 9 | Any private interior gateway (used by Cisco for IGRP) |
| udp | 17 | User Datagram |
| rdp | 27 | Reliable Data Protocol |
| ipv6 | 41 | IPv6 |
| ipv6-route | 43 | Routing Header for IPv6 |
| ipv6-frag | 44 | Fragment Header for IPv6 |
| idrp | 45 | Inter-Domain Routing Protocol |
| rsvp | 46 | Reservation Protocol |
| gre | 47 | General Routing Encapsulation |
| ipv6-icmp | 58 | ICMP for IPv6 |
| ipv6-no-nxt | 59 | No Next Header for IPv6 |
| ipv6-opts | 60 | Destination Options for IPv6 |
| iso-ip | 80 | ISO Internet Protocol |
| eigrp | 88 | EIGRP |
| ospf-igp | 89 | OSPF/IGP |
| ether-ip | 97 | Ethernet-within-IP Encapsulation |
| encap | 98 | Encapsulation Header |
| pnni | 102 | PNNI over IP |
| pim | 103 | Protocol Independent Multicast |
| vrrp | 112 | Virtual Router Redundancy Protocol |
| l2tp | 115 | Layer Two Tunneling Protocol |

Table 48 Protocol ID Descriptions (Continued)

| Protocol | Protocol ID | Description |
|----------|-------------|--------------------------------------|
| stp | 118 | Spanning Tree Protocol |
| ptp | 123 | Performance Transparency Protocol |
| isis | 124 | ISIS over IPv4 |
| crtp | 126 | Combat Radio Transport Protocol |
| crudp | 127 | Combat Radio User Datagram |
| sctp | 132 | Stream Control Transmission Protocol |

match

Syntax **match** [*next-header next-header*]
no match

Context config>filter>ipv6-filter>entry

Description This command enables the context to enter match criteria for the filter entry. When the match criteria have been satisfied the action associated with the match criteria is executed.

A **match** context may consist of multiple match criteria, but multiple **match** statements cannot be entered per entry. More precisely, the command can be entered multiple times but this only results in modifying the *next-header*, and does not affect the underlying match criteria configuration.

The **no** form of the command removes all the match criteria from the filter entry and sets the *next-header* of the match command to **none** (keyword). As per above, **match next-header none** is however not equivalent to **no match**.

Default match protocol none

Parameters *next-header* — Specifies the IPv6 next header to match. This parameter is analogous to the protocol parameter used in IPv4 filter match command.

Default Value: none (keyword)

Values [1 to 42 | 45 to 49 | 52 to 59 | 61 to 255] — in decimal format. Values can also be specified in hexadecimal format, in binary format, or using the following keywords:

none, icmp, igmp, ip, tcp, egp, igp, udp, rdp, ipv6, idrp, rsvp, gre, ipv6-icmp, ipv6-no-nxt, iso-ip, eigrp, ospf-igp, ether-ip, encap, pnni, pim, vrrp, l2tp, stp, ptp, isis, crtp, crudp, sctp

* — udp/tcp wildcard

dscp

| | |
|--------------------|--|
| Syntax | dscp <i>dscp-name</i> no dscp |
| Context | config>filter>ip-filter>entry>match config>filter>ipv6-filter>entry>match |
| Description | This command configures a DiffServ Code Point (DSCP) name to be used as an IP filter match criterion. The no form of the command removes the DSCP match criterion. |
| Default | no dscp |
| Parameters | <i>dscp-name</i> — Configures a DSCP name. The DiffServ code point may only be specified by its name. Values be, ef, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cp9, cs1, cs2, cs3, cs4, cs5, nc1, nc2, af11, af12, af13, af21, af22, af23, af31, af32, af33, af41, af42, af43, cp11, cp13, cp15, cp17, cp19, cp21, cp23, cp25, cp27, cp29, cp31, cp33, cp35, cp37, cp39, cp41, cp42, cp43, cp44, cp45, cp47, cp49, cp50, cp51, cp52, cp53, cp54, cp55, cp57, cp58, cp59, cp60, cp61, cp62, cp63 |

ah-ext-hdr

| | |
|--------------------|---|
| Syntax | ah-ext-hdr { true false } no ah-ext-hdr |
| Context | config>filter>ipv6-filter>entry>match |
| Description | This command enables match on existence of AH Extension Header in the IPv6 filter policy. The no form of this command ignores AH Extension Header presence/absence in a packet when evaluating match criteria of a given filter policy entry. |
| Default | no ah-ext-hdr |
| Parameters | true — Matches a packet with an AH Extension Header. false — Matches a packet without an AH Extension Header. |

dst-ip

| | |
|---------------|---|
| Syntax | IP-exception: dst-ip { <i>ip-address/mask</i> <i>ip-address ipv4-address-mask</i> } no dst-ip IPv4: dst-ip { <i>ip-address/mask</i> <i>ip-address ipv4-address-mask</i> ip-prefix-list <i>prefix-list</i> } |
|---------------|---|

| | |
|-------------|--|
| | <pre>name}} no dst-ip IPv6: dst-ip {ipv6-address/prefix-length ipv6-address ipv6-address-mask ipv6-prefix-list prefix-list-name} no dst-ip</pre> |
| Context | config>filter>ip-exception>entry>match config>filter>ip-filter>entry>match config>filter>ipv6-filter>entry>match |
| Description | <p>This command configures a destination address range to be used as a filter policy match criterion.</p> <p>To match on the destination address, specify the address and its associated mask, e.g., 10.1.0.0/16. The conventional notation of 10.1.0.0 255.255.0.0 can also be used for IPv4.</p> <p>The no form of this command removes the destination IPv4 or IPv6 address match criterion.</p> |
| Default | no dst-ip |
| Parameters | <p><i>ip-address</i> — Specifies the destination IPv4 address in dotted decimal notation.</p> <p>Values a.b.c.d</p> <p><i>mask</i> — Specifies the length in bits of the subnet mask.</p> <p>Values 1 to 32</p> <p><i>ipv4-address-mask</i> — Specifies the subnet mask in dotted decimal notation.</p> <p>Values a.b.c.d (dotted quad equivalent of mask length)</p> <p>ip-prefix-list or ipv6-prefix-list <i>prefix-list-name</i> — Specifies to use a list of IP prefixes referred to by <i>prefix-list-name</i>, which is a string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.</p> <p><i>ipv6-address</i> — the IPv6 prefix for the IP match criterion in hex digits.</p> <p>Values x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0..FFFF]H d: [0..255]D</p> <p><i>prefix-length</i> — the IPv6 prefix length for the specified <i>ipv6-address</i> expressed as a decimal integer.</p> <p>Values 1 to 128</p> <p><i>ipv6-address-mask</i> — eight 16-bit hexadecimal pieces representing bit match criteria.</p> <p>Values x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0..FFFF]H d: [0..255]D</p> |

dst-port

| | |
|--------------------|--|
| Syntax | dst-port { lt gt eq } <i>dst-port-number</i> dst-port port-list <i>port-list-name</i> dst-port range <i>dst-port-number dst-port-number</i> no dst-port |
| Context | config>filter>ip-exception>entry>match config>filter>ip-filter>entry>match config>filter>ipv6-filter>entry>match |
| Description | <p>This command configures a destination TCP, UDP, or SCTP port number or port range for an IP filter or IP exception match criterion. An entry containing Layer 4 non-zero match criteria will not match non-initial (2nd, 3rd, etc) fragments of a fragmented packet since only the first fragment contains the Layer 4 information. Similarly an entry containing "dst-port eq 0" match criterion, may match non-initial fragments when the destination port value is not present in a packet fragment and other match criteria are also met.</p> <p>The no form of the command removes the destination port match criterion.</p> |
| Default | no dst-port |
| Parameters | <p>lt gt eq — Specifies the operator to use relative to the <i>dst-port-number</i> for specifying the port number match criteria.</p> <p>lt specifies that all port numbers less than the <i>dst-port-number</i> match.</p> <p>gt specifies that all port numbers greater than the <i>dst-port-number</i> match.</p> <p>eq specifies that the <i>dst-port-number</i> must be an exact match.</p> <p><i>dst-port-number</i> — Specifies the destination port number to be used as a match criteria expressed as a decimal integer, as well as in hexadecimal or binary format. The following value is for decimal integer format only.</p> <p>Values 0 to 65535</p> <p><i>port-list-name</i> — Specifies to use a list of ports referred to by <i>port-list-name</i>, which is a string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.</p> <p><i>dst-port-number dst-port-number</i> — Specifies inclusive port range between two <i>dst-port-number</i> values</p> |

esp-ext-hdr

| | |
|--------------------|--|
| Syntax | esp-ext-hdr { true false } no esp-ext-hdr |
| Context | config>filter>ipv6-filter>entry>match |
| Description | This command enables match on existence of ESP Extension Header in the IPv6 filter policy. |

The **no** form of this command ignores ESP Extension Header presence/absence in a packet when evaluating match criteria of a given filter policy entry.

| | |
|-------------------|--|
| Default | no esp-ext-hdr |
| Parameters | true — Matches a packet with an ESP Extension Header. false — Matches a packet without an ESP Extension Header. |

fragment

| | |
|--------------------|--|
| Syntax | IPv4: fragment {true false} no fragment IPv6: fragment {true false first-only non-first-only} no fragment |
| Context | config>filter>ip-filter>entry>match config>filter>ipv6-filter>entry>match |
| Description | This command specifies match criterion for fragmented packets. The no form of the command removes the match criterion. |
| Default | no fragment |
| Parameters | true — Specifies to match on all fragmented IP packets false — Specifies to match on all non-fragmented IP packets first-only — matches if a packet is an initial fragment of a fragmented IPv6 packet non-first-only — matches if a packet is a non-initial fragment of a fragmented IPv6 packet |

flow-label

| | |
|--------------------|--|
| Syntax | flow-label <i>flow-label</i> [<i>mask</i>] no flow-label |
| Context | config>filter>ipv6-filter>entry>match |
| Description | This command configures the flow-label and optional mask match condition. The no form of the command reverts to the default. |
| Default | no flow-label |

| | |
|-------------------|--|
| Parameters | <i>flow-label</i> — Specifies the flow label to be used as a match criterion. Value can be expressed as a decimal integer, as well as in hexadecimal or binary format. The following value shows decimal integer format only. Values 0 to 1048575 |
| | <i>mask</i> — Specifies the flow label mask value for this policy IPv6 Filter entry. Value can be expressed as a decimal integer, as well as in hexadecimal or binary format. The following value shows decimal integer format only. Values 0 to 1048575 |

hop-by-hop-opt

| | |
|--------------------|--|
| Syntax | hop-by-hop-opt {true false} no hop-by-hop-opt |
| Context | config>filter>ipv6-filter>entry>match |
| Description | <p>This command enables match on existence of Hop-by-Hop Options Extension Header in the IPv6 filter policy.</p> <p>The no form of this command ignores Hop-by-Hop Options Extension Header presence/absence in a packet when evaluating match criteria of a given filter policy entry.</p> |
| Default | no hop-by-hop-opt |
| Parameters | true — Matches a packet with a Hop-by-Hop Options Extension header. false — Matches a packet without a Hop-by-Hop Options Extension header. |

icmp-code

| | |
|--------------------|--|
| Syntax | icmp-code <i>icmp-code</i> no icmp-code |
| Context | config>filter>ip-exception>entry>match config>filter>ip-filter>entry>match config>filter>ipv6-filter>entry>match |
| Description | <p>Configures matching on ICMP/ICMPv6 code field in the ICMP/ICMPv6 header of an IPv4 or IPv6 packet as a filter match criterion or configures matching on the ICMP code field in the ICMP header of an IPv4 packet as an exception filter match criterion. An entry containing Layer 4 non-zero match criteria will not match non-initial (2nd, 3rd, etc.) fragments of a fragmented packet since only the first fragment contains the Layer 4 information. Similarly an entry containing "icmp-code 0" match criterion, may match non-initial fragments when the Layer 4 header is not present in a packet fragment and other match criteria are also met.</p> <p>The no form of the command removes the criterion from the match entry.</p> |

| | |
|-------------------|---|
| Default | no icmp-code |
| Parameters | <i>icmp-code</i> — Specifies the ICMP/ICMPv6 code value that must be present to match. Value can be expressed as a decimal integer, as well as in hexadecimal or binary format, or even using keywords. The following value shows decimal integer only. |
| Values | 0 to 255 |

icmp-type

| | |
|--------------------|---|
| Syntax | icmp-type <i>icmp-type</i> no icmp-type |
| Context | config>filter>ip-exception>entry>match config>filter>ip-filter>entry>match config>filter>ipv6-filter>entry>match |
| Description | <p>This command configures matching on the ICMP/ICMPv6 type field in the ICMP/ICMPv6 header of an IPv4 or IPv6 packet as a filter match criterion or configures matching on the ICMP type field in the ICMP header of an IPv4 packet as an exception filter match criterion. An entry containing Layer 4 non-zero match criteria will not match non-initial (2nd, 3rd, etc.) fragments of a fragmented packet since only the first fragment contains the Layer 4 information. Similarly an entry containing "icmp-type 0" match criterion, may match non-initial fragments when the Layer 4 header is not present in a packet fragment and other match criteria are also met.</p> <p>The no form of the command removes the criterion from the match entry.</p> |
| Default | no icmp-type |
| Parameters | <i>icmp-type</i> — Specifies the ICMP/ICMPv6 type value that must be present to match. Value can be expressed as a decimal integer, as well as in hexadecimal or binary format, or even using keywords. The following value shows decimal integer only. |
| Values | 0 to 255 |

ip-option

| | |
|--------------------|--|
| Syntax | ip-option <i>ip-option-value</i> [<i>ip-option-mask</i>] no ip-option |
| Context | config>filter>ip-filter>entry>match |
| Description | <p>This command configures matching packets with a specific IP option or a range of IP options in the first option of the IP header as an IP filter match criterion.</p> <p>The option-type octet contains 3 fields:</p> <ul style="list-style-type: none"> 1 bit copied flag (copy options in all fragments) |

2 bits option class

5 bits option number

The **no** form of the command removes the match criterion.

Default no ip-option

Parameters *ip-option-value* — Specifies the 8 bit option-type as a decimal integer, binary, or hexadecimal format. The mask is applied as an AND to the option byte, the result is compared with the option-value.

The decimal value entered for the match should be a combined value of the eight bit option type field and not just the option number. Therefore, to match on IP packets that contain the Router Alert option (option number = 20), enter the option type of 148 (10010100).

Values 0 to 255

ip-option-mask — Specifies an optional parameter that can be used when specifying a range of option numbers to use as the match criteria

This 8 bit mask can be configured using the following formats:

Table 49 ip-option-mask Formats

| Format Style | Format Syntax | Example |
|--------------|---------------|-----------|
| Decimal | DDD | 20 |
| Hexadecimal | 0xHH | 0x14 |
| Binary | 0BBBBBBBB | 0b0010100 |

Default 255 (decimal) (exact match)

Values 1 to 255 (decimal)

multiple-option

Syntax **multiple-option {true | false}**
no multiple-option

Context config>filter>ip-filter>entry>match

Description This command configures matching packets that contain one or more than one option fields in the IP header as an IP filter match criterion.

The **no** form of the command removes the checking of the number of option fields in the IP header as a match criterion.

Default no multiple-option

- Parameters** **true** — Specifies matching on IP packets that contain more than one option field in the header.
- false** — Specifies matching on IP packets that do not contain multiple option fields present in the header.

option-present

- Syntax** **option-present {true | false}**
 no option-present
- Context** config>filter>ip-filter>entry>match
- Description** This command configures matching packets that contain any IP options in the IP header as an IP filter match criterion.
- The **no** form of the command removes the checking of IP options in the IP header as a match criterion.
- Default** no option-present
- Parameters** **true** — Specifies matching on all IP packets that contain any IP options in the IP header. A match will occur for all packets that have any IP option present. An option field of zero is considered as no option present.
- false** — Specifies matching on IP packets that do not have any IP option present in the IP header. (an option field of zero). An option field of zero is considered as no option present.

port

- Syntax** **port {lt | gt | eq} port-number**
 port port-list port-list-name
 port range port-number port-number
 no port
- Context** config>filter>ip-filter>entry>match
 config>filter>ipv6-filter>entry>match
- Description** This command configures a TCP/UDP/SCTP source or destination port match criterion in IPv4 and IPv6 CPM (SCTP not supported) and/or ACL filter policies. A packet matches this criterion if the packet TCP/UDP/SCTP (as configured by protocol/next-header match) source OR destination port matches either the specified port value or a port in the specified port range or port-list. .

Operational Note: This command is mutually exclusive with `src-port` and `dst-port` commands. Configuring "port eq 0", may match non-initial fragments where the source/destination port values are not present in a packet fragment if other match criteria are also met.

The **no** form of this command deletes the specified port match criterion.

| | |
|-------------------|---|
| Default | no port |
| Parameters | <p>lt gt eq — Specifies the operator to use relative to <i>port-number</i> for specifying the port number match criteria.</p> <p>lt — Specifies that all port numbers less than <i>port-number</i> match.</p> <p>gt — Specifies that all port numbers greater than <i>port-number</i> match.</p> <p>eq — Specifies that the <i>port-number</i> must be an exact match.</p> <p><i>port-number</i> — Specifies a source or destination port to be used as a match criterion. The port number can be expressed as a decimal integer, as well as in hexadecimal or binary format. The following value shows a decimal integer only.</p> <p>Values 0 to 65535</p> <p>port-list port-list-name — Specifies an inclusive range of source or destination port values to be used as match criteria.</p> <p>range port-number port-number — Specifies an inclusive range of source or destination port values to be used as match criteria.</p> |

routing-type0

| | |
|--------------------|--|
| Syntax | routing-type0 {true false} no routing-type0 |
| Context | config>filter>ipv6-filter>entry>match |
| Description | <p>This command enables match on existence of Routing Type Extension Header type 0 in the IPv6 filter policy.</p> <p>The no form of this command ignores Routing Type Extension Header type 0 presence/absence in a packet when evaluating match criteria of a given filter policy entry.</p> |
| Default | no routing-type0 |
| Parameters | <p>true — Specifies whether a packet contains Routing Type Extension Header type 0.</p> <p>false — Specifies whether a packet does not contain Routing Type Extension Header type 0 .</p> |

src-ip

| | |
|--------------------|--|
| Syntax | src-ip { <i>ip-address/mask</i> <i>ip-address ipv4-address-mask</i> } src-ip { <i>ip-address/mask</i> <i>ip-address ipv4-address-mask</i> ip-prefix-list <i>prefix-list-name</i> } src-ip { <i>ipv6-address/prefix-length</i> <i>ipv6-address ipv6-address-mask</i> ipv6-prefix-list <i>prefix-list-name</i> } no src-ip |
| Context | config>filter>ip-exception>entry>match config>filter>ip-filter>entry>match config>filter>ipv6-filter>entry>match |
| Description | <p>This command configures a source IPv4 or IPv6 address range to be used as an IP filter or IP exception match criterion.</p> <p>To match on the source IPv4 or IPv6 address, specify the address and its associated mask, for example, 10.1.0.0/16 for IPv4. The conventional notation of 10.1.0.0 255.255.0.0 may also be used for IPv4.</p> <p>The no form of the command removes the source IP address match criterion.</p> |
| Default | no src-ip |
| Parameters | <p><i>ip-address</i> — Specifies the destination IPv4 address specified in dotted decimal notation.</p> <p>Values a.b.c.d</p> <p><i>mask</i> — Specifies the length in bits of the subnet mask.</p> <p>Values 1 to 32</p> <p><i>ipv4-address-mask</i> — Specifies the subnet mask in dotted decimal notation.</p> <p>Values a.b.c.d (dotted quad equivalent of mask length)</p> <p>ip-prefix-list or ipv6-prefix-list <i>prefix-list-name</i> — Specifies to use a list of IP prefixes referred to by prefix-list-name, which is a string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.</p> <p><i>ipv6-address</i> — Specifies an IPv6 prefix for the IP match criterion in hex digits.</p> <p>Values x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:d.d.d.d x: [0..FFFF]H d: [0..255]D</p> <p><i>prefix-length</i> — Specifies whether a the IPv6 prefix length for the specified <i>ipv6-address</i> expressed as a decimal integer.</p> <p>Values 1 to 128</p> |

ipv6-address-mask — Specifies eight 16-bit hexadecimal pieces representing bit match criteria.

Values x:x:x:x:x:x:x:x (eight 16-bit pieces)
 x:x:x:x:x:x:d.d.d.d
 x: [0..FFFF]H
 d: [0..255]D

src-port

| | |
|--------------------|---|
| Syntax | src-port {lt gt eq} <i>src-port-number</i> src-port port-list <i>port-list-name</i> src-port range <i>src-port-number src-port-number</i> no src-port |
| Context | config>filter>ip-exception>entry>match config>filter>ip-filter>entry>match config>filter>ipv6-filter>entry>match |
| Description | <p>This command configures a source TCP, UDP, or SCTP port number, port range, or port match list for an IP filter or IP exception match criterion. An entry containing Layer 4 non-zero match criteria will not match non-initial (2nd, 3rd, and so on) fragments of a fragmented packet since only the first fragment contains the Layer 4 information. Similarly an entry containing "src-port eq 0" match criterion, may match non-initial fragments when the source port value is not present in a packet fragment and other match criteria are also met.</p> <p>The no form of the command removes the source port match criterion.</p> |
| Default | no src-port |
| Parameters | <p>lt gt eq — Specifies the operator to use relative to <i>src-port-number</i> for specifying the port number match criteria.</p> <p>lt specifies that all port numbers less than <i>src-port-number</i> match.</p> <p>gt specifies that all port numbers greater than <i>src-port-number</i> match.</p> <p>eq specifies that <i>src-port-number</i> must be an exact match.</p> <p><i>src-port-number</i> — Specifies the source port number to be used as a match criteria expressed as a decimal integer, and in hexadecimal or binary format. Below shows decimal integer only.</p> <p>Values 0 to 65535</p> <p><i>port-list-name</i> — Specifies to use a list of ports referred to by <i>port-list-name</i>, which is a string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes.</p> <p><i>src-port-number src-port-number</i> — Specifies inclusive port range between two src-port-number values</p> |

src-route-option

| | |
|--------------------|--|
| Syntax | src-route-option {true false} no source-route-option |
| Context | config>filter>ip-filter>entry>match |
| Description | This command enables source route option match conditions. When enabled, this filter should match if a (strict or loose) source route option is present/not present at any location within the IP header, as per the value of this object. The no form of the command removes the criterion from the match entry. |
| Default | no src-route-option |
| Parameters | true — Enables source route option match conditions. false — Disables source route option match conditions. |

tcp-ack

| | |
|--------------------|---|
| Syntax | tcp-ack {true false} no tcp-ack |
| Context | config>filter>ip-filter>entry>match config>filter>ipv6-filter>entry>match |
| Description | This command configures matching on the ACK bit being set or reset in the control bits of the TCP header of an IP packet as an IP filter match criterion. An entry containing Layer 4 non-zero match criteria will not match non-initial (2nd, 3rd, etc.) fragments of a fragmented packet since only the first fragment contains the Layer 4 information. The no form of the command removes the criterion from the match entry. |
| Default | no tcp-ack |
| Parameters | true — Specifies matching on IP packets that have the ACK bit set in the control bits of the TCP header of an IP packet. false — Specifies matching on IP packets that do not have the ACK bit set in the control bits of the TCP header of the IP packet. |

tcp-syn

| | |
|----------------|--|
| Syntax | tcp-syn {true false} no tcp-syn |
| Context | config>filter>ip-filter>entry>match config>filter>ipv6-filter>entry>match |

| | |
|--------------------|--|
| Description | <p>This command configures matching on the SYN bit being set or reset in the control bits of the TCP header of an IP packet as an IP filter match criterion. An entry containing Layer 4 non-zero match criteria will not match non-initial (2nd, 3rd, etc.) fragments of a fragmented packet since only the first fragment contains the Layer 4 information.</p> <p>The SYN bit is normally set when the source of the packet wants to initiate a TCP session with the specified destination IP address.</p> <p>The no form of the command removes the criterion from the match entry.</p> |
| Default | no tcp-syn |
| Parameters | <p>true — Specifies matching on IP packets that have the SYN bit set in the control bits of the TCP header.</p> <p>false — Specifies matching on IP packets that do not have the SYN bit set in the control bits of the TCP header.</p> |

4.4.2.7 Match List Configuration Commands

match-list

| | |
|--------------------|--|
| Syntax | match-list |
| Context | config>filter |
| Description | This command enables the configuration context for match lists to be used in filter policies (IOM/FP and CPM). |
| Default | n/a |

ip-prefix-list

| | |
|--------------------|---|
| Syntax | ip-prefix-list <i>ip-prefix-list-name</i> [create] no ip-prefix-list <i>ip-prefix-list-name</i> |
| Context | config>filter>match-list |
| Description | <p>This command creates a list of IPv4 prefixes for match criteria in IPv4 ACL and CPM filter policies.</p> <p>The no form of this command deletes the specified list.</p> <p>Operational Notes:</p> <p>An ip-prefix-list must contain only IPv4 address prefixes.</p> |

An IPv4 prefix match list cannot be deleted if it is referenced by a filter policy.

See general description related to match-list usage in filter policies.

| | |
|-------------------|---|
| Default | n/a |
| Parameters | <i>ip-prefix-list-name</i> — Specifies a string of up to 32 printable ASCII characters. If special characters are used, the string must be enclosed within double quotes. |

ipv6-prefix-list

| | |
|--------------------|--|
| Syntax | ipv6-prefix-list <i>ipv6-prefix-list-name</i> [create] no ipv6-prefix-list <i>ipv6-prefix-list-name</i> |
| Context | config>filter>match-list |
| Description | This command creates a list of IPv6 prefixes for match criteria in ACL and CPM IPv6 filter policies. The no form of this command deletes the specified list. Operational Notes: An ipv6-prefix-list must contain only IPv6 address prefixes. An ipv6-prefix-list cannot be deleted if it is referenced by a filter policy. See general description related to match-list usage in filter policies. |
| Default | n/a |
| Parameters | <i>ipv6-prefix-list-name</i> — Specifies a string of up to 32 printable ASCII characters. If special characters are used, the string must be enclosed within double quotes. |

apply-path

| | |
|--------------------|--|
| Syntax | [no] apply-path no apply-path |
| Context | config>filter>match-list>ip-prefix-list config>filter>match-list>ipv6-prefix-list |
| Description | This command enables the context to configure auto-generation of address prefixes for IPv4 or IPv6 address prefix match lists. The context in which the command is executed governs whether IPv4 or IPv6 prefixes will be auto-generated. The no form of this command removes all auto-generation configuration under the apply-path context. |

Default no apply path

bgp-peers

| | |
|--------------------|---|
| Syntax | bgp-peers <i>criterion-index</i> group <i>reg-exp</i> neighbor <i>reg-exp</i> bgp-peers <i>criterion-index</i> router <i>router-instance</i> group <i>reg-exp</i> neighbor <i>reg-exp</i> bgp-peers <i>criterion-index</i> router service-name <i>service-name</i> group <i>reg-exp</i> neighbor <i>reg-exp</i> no bgp-peers <i>criterion-index</i> |
| Context | config>filter>match-list>ip-prefix-list>apply-path config>filter>match-list>ipv6-prefix-list>apply-path |
| Description | <p>This command configures auto-generation of IPv4 or IPv6 address prefixes (as required by the context that the command is executed within) based on the base router BGP instance configuration.</p> <p>The no form of this command removes the bgp-peers configuration for auto-generation of address prefixes for the specified index value.</p> |
| Default | No embedded filter policies are included in a filter policy. |
| Parameters | <p><i>service-name</i> — Specifies the service name, up to 64 characters in length.</p> <p>group — Configures a match against the base router BGP instance group configuration. Regex wildcard match (.) can be used to match against any group.</p> <p>neighbor — Configures a match against the base router BGP instance neighbor configuration. Regex wildcard match (.) can be used to match against any neighbor.</p> <p><i>criterion-index</i> — Specifies an integer from 1 to 255 enumerating BGP peers auto-generation configuration within this list.</p> <p><i>router-instance</i> — Specifies the router name or service ID.</p> <p>Values</p> <ul style="list-style-type: none"> router-instance: <i>router-name</i> or <i>vprn-svc-id</i> router-name: "Base" vprn-svc-id: 1 to 2147483647 <i>service-name</i>: Specifies the service name, up to 64 characters in length. <p>router — Configures a match against the base router BGP instance configuration.</p> <p><i>reg-exp</i> — Specifies a regular expression that defines a match string, up to 255 characters in length, to be used to auto-generate address prefixes. Matching is performed from the least-significant digit. For example, a string 10.0 matches all neighbors with addresses starting with 10, such as 10.0.x.x or 10.0xx.x.x.</p> |

port-list

| | |
|--------------------|--|
| Syntax | port-list <i>port-list-name</i> [create] no port-list <i>port-list-name</i> |
| Context | config>filter>match-list |
| Description | This command creates a list of TCP/UDP/SCTP port values or ranges for match criteria in IPv4 and IPv6 ACL and CPM filter policies. The no form of this command deletes the specified list. Operational notes: SCTP port match is supported in ACL filter policies only. A port-list must contain only TCP/UDP/SCTP port values or ranges. A TCP/UDP/SCTP port match list cannot be deleted if it is referenced by a filter policy. See general description related to match-list usage in filter policies. |
| Parameters | <i>port-list-name</i> — Specifies a string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes. |

port

| | |
|--------------------|--|
| Syntax | [no] port <i>port-number</i> [no] port range <i>start end</i> |
| Context | config>filter>match-list>port-list |
| Description | This command adds a port or a range of ports to an existing port match list. The no form of this command deletes the specified port or range of ports from the list. |
| Parameters | <i>port-number</i> — Specifies the port number to add to the list. The port number can be expressed as a decimal integer, as well as in hexadecimal or binary format. Below shows decimal integer only. Values 0 to 65535 <i>start end</i> — Specifies an inclusive port range between two port numbers values. The <i>start</i> of the range and <i>end</i> of the range can be expressed as decimal integers, as well as in hexadecimal or binary format. The following value shows decimal integer only. Values 0 to 65535 |

prefix

| | |
|--------------------|---|
| Syntax | [no] prefix <i>ipv6-prefix/prefix-length</i> |
| Context | config>filter>match-list>ipv6-prefix-list |
| Description | <p>This command adds an IPv6 address prefix to an existing IPv6 address prefix match list.</p> <p>The no form of this command deletes the specified prefix from the list.</p> <p>Operational Notes:</p> <p>To add set of different prefixes, execute the command with all unique prefixes. The prefixes are allowed to overlap IPv6 address space.</p> <p>An IPv6 prefix addition will be blocked, if resource exhaustion is detected anywhere in the system because of filter policies that use this IPv6 address prefix list.</p> |
| Parameters | <p><i>ipv6-prefix</i> — Specifies an IPv6 address prefix written as hexadecimal numbers separated by colons with host bits set to 0. One string of zeros can be omitted so 1010::700:0:217A is equivalent to 1010:0:0:0:700:0:217A.</p> <p>Values x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x: [0..FFFF]H d: [0..255]D</p> <p><i>prefix-length</i> — Specifies the length of the entered IPv6 prefix.</p> <p>Values 1 to 128</p> |

prefix

| | |
|--------------------|--|
| Syntax | [no] prefix <i>ip-prefix/prefix-length</i> |
| Context | config>filter>match-list>ip-prefix-list |
| Description | <p>This command adds an IPv4 address prefix to an existing IPv4 address prefix match list.</p> <p>The no form of this command deletes the specified prefix from the list.</p> <p>Operational Notes:</p> <p>To add set of unique prefixes, execute the command with all unique prefixes. The prefixes are allowed to overlap IPv4 address space.</p> <p>An IPv4 prefix addition will be blocked, if resource exhaustion is detected anywhere in the system because of filter policies that use this IPv4 address prefix list.</p> |

| | |
|-------------------|--|
| Parameters | <i>ip-prefix</i> — Specifies a valid IPv4 address prefix in dotted decimal notation. |
| Values | 0.0.0.0 to 255.255.255.255 (host bit must be 0) |
| | <i>prefix-length</i> — Specifies the length of the entered IPv4 prefix. |
| Values | 0 to 32 |

4.4.2.8 MAC Filter Entry Commands

action

| | |
|--------------------|---|
| Syntax | [no] action [secondary] |
| Context | config>filter>mac-filter>entry |
| Description | This command sets the context for specific action commands to be performed on packets matching this filter entry. |

drop

| | |
|--------------------|--|
| Syntax | drop |
| Context | config>filter>mac-filter>entry>action |
| Description | This command sets the MAC filter entry action to drop. |

forward

| | |
|--------------------|---|
| Syntax | forward forward esi esi service-id vpls-service-id forward sap sap-id forward sdp sdp-id:vc-id |
| Context | config>filter>mac-filter>entry>action |
| Description | This command sets the context for specific forward commands to be performed. |

match

| | |
|---------------|---|
| Syntax | match [frame-type {802dot3 802dot2-llc 802dot2-snap ethernet_II}] no match |
|---------------|---|

| | |
|--------------------|---|
| Context | config>filter>mac-filter>entry |
| Description | <p>This command creates the context for entering/editing match criteria for the filter entry and specifies an Ethernet frame type for the entry.</p> <p>A match context may consist of multiple match criteria, but multiple match statements cannot be entered per entry.</p> <p>The no form of the command removes the match criteria for the <i>entry-id</i>.</p> |
| Default | n/a |
| Parameters | <p>frame-type <i>keyword</i> — the frame-type keyword configures an Ethernet frame type to be used for the MAC filter match criteria</p> <p>Default 802dot3</p> <p>Values 802dot3, 802dot2-llc, 802dot2-snap, ethernet_II</p> <p>802dot3 — Specifies the frame type is Ethernet IEEE 802.3</p> <p>802dot2-llc — Specifies the frame type is Ethernet IEEE 802.2 LLC</p> <p>802dot2-snap — Specifies the frame type is Ethernet IEEE 802.2 SNAP</p> <p>ethernet_II — Specifies the frame type is Ethernet Type II</p> |

4.4.2.9 MAC Filter Match Criteria

rate-limit

| | |
|--------------------|--|
| Syntax | rate-limit <i>value</i> |
| Context | config>filter>mac-filter>entry>action |
| Description | <p>This command sets the rate limit for the traffic matching both the filter entry match criteria and the <i>packet-length-value</i> defined in the rate-limit action statement.</p> <p>Packets matching the filter entry match criteria and not matching the <i>packet-length-value</i> defined in the rate-limit action statement are implicitly forwarded with no further match in subsequent filter entries.</p> <p>Rate limit packets matching both the filter entry match criteria and the <i>tll-value</i> are defined in the action rate-limit statement.</p> <p>Packets matching the filter entry match criteria and not matching the <i>tll-value</i> defined in the rate-limit action statement are implicitly forwarded with no further match in following filter entries.</p> |

Parameters *value* — Specifies the **rate-limit** *value* in kb/s. A rate of 0 results in all traffic being dropped. A rate of **max** results in all traffic being forwarded.

Values 0 to 2000000000 kb/s | max

dot1p

Syntax **dot1p** *dot1p-value* [*dot1p-mask*]
no dot1p

Context config>filter>mac-filter>entry>match

Description Configures an IEEE 802.1p value or range to be used as a MAC filter match criterion.

When a frame is missing the 802.1p bits, specifying an dot1p match criterion will fail for the frame and result in a non-match for the MAC filter entry.

The **no** form of the command removes the criterion from the match entry.

Egress **dot1p** value matching will only match if the customer payload contains the 802.1p bits. For example, if a packet ingresses on a null encapsulated SAP and the customer packet is IEEE 802.1Q or 802.1p tagged, the 802.1p bits will be present for a match evaluation. On the other hand, if a customer tagged frame is received on a dot1p encapsulated SAP, the tag will be stripped on ingress and there will be no 802.1p bits for a MAC filter match evaluation; in this case, any filter entry with a dot1p match criterion specified will fail.

Default no dot1p

Parameters *dot1p-value* — Specifies the IEEE 802.1p value in decimal.

Values 0 to 7

dot1p-mask — Specifies a 3-bit mask that can be configured using the decimal integer, hexadecimal or binary format.

Table 50 dot1p-mask Formats

| Format Style | Format Syntax | Example |
|--------------|---------------|---------|
| Decimal | D | 4 |
| Hexadecimal | 0xH | 0x4 |
| Binary | 0bBBB | 0b100 |

To select a range from 4 up to 7 specify *dot1p-value* of 4 and a *dot1p-mask* of 0b100 for value and mask.

Default 7 (decimal)
Values 1 to 7 (decimal)

dsap

- Syntax** **dsap** *dsap-value* [*dsap-mask*]
no dsap
- Context** config>filter>mac-filter>entry>match
- Description** Configures an Ethernet 802.2 LLC DSAP value or range for a MAC filter match criterion.
- This is a one-byte field that is part of the 802.2 LLC header of the IEEE 802.3 Ethernet Frame.
- The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and may not be part of the same match criteria.
- Use the **no** form of the command to remove the dsap value as the match criterion.
- Default** no dsap
- Parameters** *dsap-value* — Specifies the 8-bit dsap match criteria value which can be expressed in decimal integer, hexadecimal or binary format.
- Values** 0 to 255
- dsap-mask* — Specifies an optional parameter that may be used when specifying a range of dsap values to use as the match criteria.
- This 8 bit mask can be configured using the decimal integer, hexadecimal or binary formats described in [Table 51](#).

Table 51 dsap-mask Formats

| Format Style | Format Syntax | Example |
|--------------|---------------|------------|
| Decimal | DDD | 240 |
| Hexadecimal | 0xHH | 0xF0 |
| Binary | 0bBBBBBBBB | 0b11110000 |

- Default** **255 (exact match)**
0x00 to 0xFF
- Values** 0 to 255

dst-mac

- Syntax** **dst-mac** *ieee-address* [*ieee-address-mask*]
no dst-mac
- Context** config>filter>mac-filter>entry>match

| | |
|--------------------|---|
| Description | Configures a destination MAC address or range to be used as a MAC filter match criterion. The no form of the command removes the destination mac address as the match criterion. |
| Default | no dst-mac |
| Parameters | <i>ieee-address</i> — Specifies the MAC address to be used as a match criterion. Values HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit. Note that both upper and lower case are supported. <i>ieee-address-mask</i> — Specifies a 48-bit mask to match a range of MAC address values. To configure so that all packets with a destination MAC OUI value of 00:03:FA are subject to a match condition then the entry should be specified as: 00:03:FA:00:00:00 FF:FF:FF:00:00:00. Default ff:ff:ff:ff:ff:ff (exact match) Values HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit. to 0xFFFFFFFFFFFF Note that both upper and lower case are supported. |

etype

| | |
|--------------------|---|
| Syntax | etype <i>0x0600..0xffff</i> no etype |
| Context | config>filter>mac-filter>entry>match |
| Description | Configures an Ethernet type II Ethertype value to be used as a MAC filter match criterion. The Ethernet type field is a two-byte field used to identify the protocol carried by the Ethernet frame. For example, 0800 is used to identify the IPv4 packets. The Ethernet type field is used by the Ethernet version-II frames. IEEE 802.3 Ethernet frames do not use the type field. For IEEE 802.3 frames, use the dsap, ssap or snap-pid fields as match criteria. The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and may not be part of the same match criteria. The no form of the command removes the previously entered etype field as the match criteria. |
| Default | no etype |
| Parameters | <i>0x0600..0xffff</i> — Specifies the Ethernet type II frame Ethertype value to be used as a match criterion expressed in decimal integer or hexadecimal format. Values 1536 to 65535 or 0x0600 to 0xFFFF |

isid

| | |
|--------------------|--|
| Syntax | isid <i>value</i> [<i>to higher-value</i>] no isid |
| Context | config>filter>mac-filter>entry>match |
| Description | <p>This command configures an ISID value or a range of ISID values to be matched by the mac-filter parent. The pbb-etype value for the related SAP (inherited from the ethernet port configuration) or for the related SDP binding (inherited from SDP configuration) will be used to identify the ISID tag.</p> <p>The no form of this command removes the ISID match criterion.</p> |
| Default | no isid |
| Parameters | <p><i>value</i> — Specifies the ISID value, 24 bits as a decimal integer. When just one present identifies a specific ISID to be used for matching.</p> <p>Values 0 to 16777215</p> <p><i>to higher-value</i> — Identifies a range of ISIDs to be used as matching criteria.</p> |

inner-tag

| | |
|--------------------|---|
| Syntax | inner-tag <i>value</i> [<i>vid-mask</i>] no inner-tag |
| Context | config>filter>mac-filter>entry>match |
| Description | <p>This command configures the matching of the second tag that is carried transparently through the service. The inner-tag on ingress is the second tag on the frame if there are no service delimiting tags. Inner tag is the second tag before any service delimiting tags on egress but is dependent in the ingress configuration and may be set to 0 even in cases where additional tags are on the frame. This allows matching VLAN tags for explicit filtering or QoS setting when using default or null encapsulations.</p> <p>The inner-tag is not applicable in ingress on dot1Q SAPs. The inner-tag may be populated on egress depending on the ingress SAP type.</p> <p>On QinQ SAPs of null and default that do not strip tags inner-tag will contain the second tag (which is still the second tag carried transparently through the service.) On ingress SAPs that strip any tags, inner-tag will contain 0 even if there are more than 2 tags on the frame.</p> <p>The optional <i>vid-mask</i> is defaulted to 4095 (exact match) but may be specified to allow pattern matching. The masking operation is ((value and vid-mask) == (tag and vid-mask)). A value of 6 and a mask of 7 would match all VIDs with the lower 3 bits set to 6.</p> <p>For QoS the VID type cannot be specified on the default QoS policy.</p> <p>The default vid-mask is set to 4095 for exact match.</p> |

Default no inner-tag

outer-tag

Syntax **outer-tag** *value* [*vid-mask*]
no outer-tag

Context config>filter>mac-filter>entry>match

Description This command configures the matching of the first tag that is carried transparently through the service. Service delimiting tags are stripped from the frame and outer tag on ingress is the first tag after any service delimiting tags. Outer tag is the first tag before any service delimiting tags on egress. This allows matching VLAN tags for explicit filtering or QoS setting when using default or null encapsulations.

On dot1Q SAPs outer-tag is the only tag that can be matched. On dot1Q SAPs with exact match (sap 2/1/1:50) the outer-tag will be populated with the next tag that is carried transparently through the service or 0 if there is no additional VLAN tags on the frame.

On QinQ SAPs that strip a single service delimiting tag, outer-tag will contain the next tag (which is still the first tag carried transparently through the service.) On SAPs with two service delimiting tags (two tags stripped) outer-tag will contain 0 even if there are more than 2 tags on the frame.

The optional *vid-mask* is defaulted to 4095 (exact match) but may be specified to allow pattern matching. The masking operation is ((value & vid-mask) == (tag & vid-mask)). A value of 6 and a mask of 7 would match all VIDs with the lower 3 bits set to 6.

For QoS the VID type cannot be specified on the default QoS policy.

The default vid-mask is set to 4095 for exact match.

Default no outer-tag

snap-oui

Syntax **snap-oui** {**zero** | **non-zero**}
no snap-oui

Context config>filter>mac-filter>entry>match

Description This command configures an IEEE 802.3 LLC SNAP Ethernet Frame OUI zero or non-zero value to be used as a MAC filter match criterion.

The **no** form of the command removes the criterion from the match criteria.

Default no snap-oui

-
- Parameters** **zero** — Specifies to match packets with the three-byte OUI field in the SNAP-ID set to zero
- non-zero** — Specifies to match packets with the three-byte OUI field in the SNAP-ID not set to zero

snap-pid

- Syntax** **snap-pid** *snap-pid*
 no snap-pid
- Context** config>filter>mac-filter>entry>match
- Description** Configures an IEEE 802.3 LLC SNAP Ethernet Frame PID value to be used as a MAC filter match criterion.
- This is a two-byte protocol id that is part of the IEEE 802.3 LLC SNAP Ethernet Frame that follows the three-byte OUI field.
- The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and may not be part of the same match criteria.
- The snap-pid match criterion is independent of the OUI field within the SNAP header. Two packets with different three-byte OUI fields but the same PID field will both match the same filter entry based on a snap-pid match criteria.
- The **no** form of the command removes the snap-pid value as the match criteria.
- Default** no snap-pid
- Parameters** *snap-pid* — Specifies the two-byte snap-pid value to be used as a match criterion. The value can be expressed in decimal integer or hexadecimal format.
- Values** 0 to 65535 or 0x0000 to 0xFFFF

src-mac

- Syntax** **src-mac** *ieee-address* [*ieee-address-mask*]
 no src-mac
- Context** config>filter>mac-filter>entry>match
- Description** Configures a source MAC address or range to be used as a MAC filter match criterion.
- The **no** form of the command removes the source mac as the match criteria.
- Default** no src-mac

| | |
|-------------------|---|
| Parameters | <i>ieee-address</i> — Specifies the 48-bit IEEE mac address to be used as a match criterion |
| Values | HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is a hexadecimal digit; both upper and lower case are supported. |
| | <i>ieee-address-mask</i> — a 48-bit mask to match a range of MAC address values |
| | To configure so that all packets with a source MAC OUI value of 00:03:FA are subject to a match condition then the entry should be specified as: 00:03:FA:00:00:00 FF:FF:FF:00:00:00 |
| Default | ff:ff:ff:ff:ff:ff (exact match) |
| Values | HH:HH:HH:HH:HH:HH or HH-HH-HH-HH-HH-HH where H is an hexadecimal digit; both upper and lower case are supported. |

ssap

| | |
|--------------------|--|
| Syntax | ssap <i>ssap-value</i> [<i>ssap-mask</i>] no ssap |
| Context | config>filter>mac-filter>entry>match |
| Description | <p>This command configures an Ethernet 802.2 LLC SSAP value or range for a MAC filter match criterion.</p> <p>This is a one-byte field that is part of the 802.2 LLC header of the IEEE 802.3 Ethernet Frame.</p> <p>The snap-pid field, etype field, ssap and dsap fields are mutually exclusive and may not be part of the same match criteria.</p> <p>The no form of the command removes the ssap match criterion.</p> |
| Default | no ssap |
| Parameters | <i>ssap-value</i> — Specifies the 8-bit ssap match criteria value in decimal, hexadecimal or binary. Values 0 to 255 |
| | <i>ssap-mask</i> — Specifies an optional parameter that may be used when specifying a range of ssap values to use as the match criteria. This 8 bit mask and the ssap value can be configured as described in Table 52 . |

Table 52 8-bit Mask Syntax

| Format Style | Format Syntax | Example |
|--------------|---------------|------------|
| Decimal | DDD | 240 |
| Hexadecimal | 0xHH | 0xF0 |
| Binary | 0bBBBBBBBB | 0b11110000 |

Values 0 to 255

4.4.2.10 IP Exception Filter Policy Commands

ip-exception

Syntax **ip-exception** *filter-id* [**create**]
[**no**] **ip-exception** {*filter-id* | *filter-name*}

Context config>filter

Description This command creates a configuration context for an IPv4 exception filter policy. After creating an exception filter ID, you can optionally assign it to a unique name with the [filter-name](#) command. The exception filter name can be used instead of the ID for exception configuration commands, show commands, monitor commands, clear commands, and port and interface association commands.

IP exception filter policies specify matching criteria that allow a packet to be an exception to where it is applied. For more information, refer to the **ip-exception** command in [Router Interface Commands](#).

The IP exception filter policy is a template that can be applied to multiple router interface group encryption contexts as long as the **scope** of the policy is configured as **template**.

Any changes made to the existing policy, using any subcommands, are applied immediately to all network interfaces where the policy is applied.

The **no** form of the command deletes the IP exception filter policy. An exception filter policy cannot be deleted until it is removed from all network interfaces where it is applied.

Parameters *filter-id* — the IP exception filter policy ID number

Values 1 to 65535

filter-name — the IP exception filter policy name, up to 64 characters in length. The name must already exist within the created IP exceptions.

create — keyword required when first creating the configuration context. When the context is created, you can navigate into the context without the **create** keyword.

4.4.2.11 Policy and Entry Maintenance Commands

copy

| | |
|--------------------|--|
| Syntax | copy ip-filter <i>src-filter-id</i> [src-entry <i>src-entry-id</i>] to <i>dst-filter-id</i> [dst-entry <i>dst-entry-id</i>] [overwrite] copy ipv6-filter <i>src-filter-id</i> [src-entry <i>src-entry-id</i>] to <i>dst-filter-id</i> [dst-entry <i>dst-entry-id</i>] [overwrite] copy mac-filter <i>src-filter-id</i> [src-entry <i>src-entry-id</i>] to <i>dst-filter-id</i> [dst-entry <i>dst-entry-id</i>] [overwrite] |
| Context | config>filter |
| Description | <p>This command copies existing filter list entries for a specific filter ID to another filter ID. The copy command is a configuration level maintenance tool used to create new filters using existing filters. It also allows bulk modifications to an existing policy with the use of the overwrite keyword.</p> <p>If overwrite is not specified, an error will occur if the destination policy ID exists.</p> |
| Default | n/a |
| Parameters | <p>ip-filter — This keyword indicates that the <i>src-filter-id</i> and the <i>dst-filter-id</i> are IPv4 filter IDs.</p> <p>ipv6-filter — This keyword indicates that the <i>src-filter-id</i> and the <i>dst-filter-id</i> are IPv6 filter IDs.</p> <p>mac-filter — This keyword indicates that the <i>src-filter-id</i> and the <i>dst-filter-id</i> are MAC filter IDs.</p> <p><i>src-filter-id</i> — Identifies the source filter policy from which the copy command will attempt to copy. The filter policy must exist within the context of the preceding keyword (ip-filter, ipv6-filter or mac-filter).</p> <p><i>dst-filter-id</i> — Identifies the destination filter policy to which the copy command will attempt to copy. If the overwrite keyword does not follow, the filter policy ID cannot already exist within the system for the filter type the copy command is issued for. If the overwrite keyword is present, the destination policy ID may or may not exist.</p> <p>overwrite — This keyword specifies that the destination filter ID may exist. If it does, everything in the existing destination filter ID will be completely overwritten with the contents of the source filter ID. If the destination filter ID exists, either overwrite must be specified or an error message will be returned. If overwrite is specified, the function of copying from source to destination occurs in a 'break before make' manner and therefore should be handled with care.</p> |

renum

| | |
|--------------------|---|
| Syntax | renum <i>old-entry-id</i> <i>new-entry-id</i> |
| Context | config>filter>ip-exception config>filter>ip-filter config>filter>ipv6-filter config>filter>mac-filter |
| Description | <p>This command renumbers existing MAC, IPv4/IPv6, or IP exception filter entries to properly sequence filter entries.</p> <p>This may be required in some cases since the OS exits when the first match is found and executes the actions according to the accompanying action command. This requires that entries be sequenced correctly from most to least explicit.</p> |
| Default | n/a |
| Parameters | <p><i>old-entry-id</i> — Specifies the entry number of an existing entry, as a decimal integer.</p> <p>Values 1 to 65535</p> <p><i>new-entry-id</i> — Specifies the new entry-number to be assigned to the old entry, as a decimal integer.</p> <p>Values 1 to 65535</p> |

4.4.2.12 Redirect Policy Commands

destination

| | |
|--------------------|---|
| Syntax | destination <i>ip-address</i> [create] no destination <i>ip-address</i> |
| Context | config>filter>redirect-policy |
| Description | <p>This command defines a destination in a redirect policy. More than one destination can be configured. Whether a destination IPv4/IPv6 address will receive redirected packets depends on the effective priority value after evaluation.</p> <p>The most preferred destination is programmed in hardware as action forward next-hop. If all destinations are down (as determined by the supported tests), action forward is programmed in hardware. All destinations within a given policy must be either IPv4 or (exclusive) IPv6. The redirect policy with IPv4 destinations configured can only be used by IPv4 filter policies. The redirect policy with IPv6 destinations configured can only be used by IPv6 filter policies.</p> |
| Default | no destination |

| | | | |
|-------------------|---|-------------------------------------|--|
| Parameters | <i>ip-address</i> — Specifies the IPv4 address (in dotted decimal notation) or IPv6 address to send the redirected traffic to | | |
| Values | IPv4 address: | ip-address: a.b.c.d | |
| | IPv6-address: | x:x:x:x:x:x:x (eight 16-bit pieces) | |
| | | x:x:x:x:x:x:d.d.d.d | |
| | x: | [0..FFFF]H | |
| | d: | [0..255]D | |

sticky-dest

| | | | |
|--------------------|---|--|--|
| Syntax | sticky-dest no-hold-time-up sticky-dest hold-time-up no sticky-dest | | |
| Context | config>filter>redirect-policy | | |
| Description | <p>This command configures sticky destination behavior for redirect policy. When enabled, the active destination is not changed to a new better destination, unless the active destination goes down or manual switch is forced using the tools>perform>filter>redirect-policy>activate-best-dest command.</p> <p>The <i>hold-time-up</i> parameter allows the operator to delay programming of the PBR to the most-preferred destination for a specified amount of time when the first destination comes up (action forward remains in place). When the first destination comes up, the timer is started and upon the expiry, the current most-preferred destination is selected (which may differ from the one that triggered the timer to start) and programmed as a sticky PBR destination. Changing the value of the timer, while the timer is running takes immediate effect.</p> <p>The no form of the command disables sticky destination behavior.</p> | | |
| Default | no sticky-dest | | |
| Parameters | <i>hold-time-up</i> — Specifies the initial delay in seconds. | | |
| Values | 0 to 65535 where 0 is equivalent to no-hold-time-up | | |

ping-test

| | | | |
|--------------------|--|--|--|
| Syntax | [no] ping-test | | |
| Context | config>filter>redirect-policy>dest | | |
| Description | This command configures parameters to perform connectivity ping tests to validate the ability for the destination to receive redirected traffic. | | |

Default no ping-test

drop-count

| | |
|--------------------|--|
| Syntax | drop-count <i>consecutive-failures</i> [hold-down <i>seconds</i>] no drop-count |
| Context | config>filter>redirect-policy>dest>ping-test config>filter>redirect-policy>dest>snmp-test config>filter>redirect-policy>dest>url-test |
| Description | This command specifies the number of consecutive requests that must fail for the destination to be declared unreachable and the time to hold destination unreachable before repeating tests. |
| Default | drop-count 3 hold-down 0 |
| Parameters | <i>consecutive-failures</i> — Specifies the number of consecutive ping test failures before declaring the destination down. Values 1 to 60 <i>hold-down seconds</i> — Specifies the amount of time, in seconds, that the system should be held down if any of the test has marked it unreachable. Values 0 to 86400 |

interval

| | |
|--------------------|---|
| Syntax | interval <i>seconds</i> no interval |
| Context | config>filter>redirect-policy>dest>ping-test config>filter>redirect-policy>dest>snmp-test config>filter>redirect-policy>dest>url-test |
| Description | This command specifies the amount of time, in seconds, between consecutive requests sent to the far end host. |
| Default | interval 1 |
| Parameters | <i>seconds</i> — Specifies the amount of time, in seconds, between consecutive requests sent to the far end host. Values 1 to 60 |

timeout

| | |
|--------------------|--|
| Syntax | timeout <i>seconds</i> no timeout |
| Context | config>filter>redirect-policy>dest>ping-test config>filter>redirect-policy>dest>snmp-test config>filter>redirect-policy>dest>url-test |
| Description | Specifies the amount of time, in seconds, that is allowed for receiving a response from the far-end host. If a reply is not received within this time the far-end host is considered unresponsive. |
| Default | timeout 1 |
| Parameters | <i>seconds</i> — Specifies the amount of time, in seconds, that is allowed for receiving a response from the far end host. Values 1 to 60 |

priority

| | |
|--------------------|---|
| Syntax | priority <i>priority</i> no priority |
| Context | config>filter>redirect-policy>dest |
| Description | Redirect policies can contain multiple destinations. Each destination is assigned an initial or base priority which describes its relative importance within the policy. |
| Default | priority 100 |
| Parameters | <i>priority</i> — Specifies the priority, expressed as a decimal integer, used to weigh the destination's relative importance within the policy. Values 1 to 255 |

snmp-test

| | |
|--------------------|---|
| Syntax | [no] snmp-test <i>test-name</i> |
| Context | config>filter>redirect-policy>dest |
| Description | This command enables the context to configure SNMP test parameters. |
| Default | n/a |

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

oid

Syntax **oid** *oid-string* **community** *community-string*
no oid

Context config>filter>redirect-policy>dest>snmp-test

Description This command specifies the OID of the object to be fetched from the destination.

Default no oid

Parameters *oid-string* — Specifies the object identifier (OID) in the OID field
community-string — Specifies the SNMP v2 community string or the SNMP v3 context name used to conduct this SNMP test.

return-value

Syntax **return-value** *return-value* **type** *return-type* [**disable** | **lower-priority** *priority* | **raise-priority** *priority*]
no return-value *return-value* **type** *return-type*

Context config>filter>redirect-policy>dest>snmp-test

Description This command specifies the criterion to adjust the priority based on the test result. Multiple criteria can be specified with the condition that they are not conflicting or overlap. If the returned value is within the specified range, the priority can be disabled, lowered or raised.

Default n/a

Parameters *return-value* — Specifies the SNMP value against which the test result is matched up to 256 characters in length.
return-type — Specifies the SNMP object type against which the test result is matched.
Values integer, unsigned, string, ip-address, counter, time-ticks, opaque
disable — Specifies that the destination may not be used for the amount of time specified in the hold-time command when the test result matches the criterion.
priority — Specifies the amount to lower the priority of the destination.
Values 1 to 255
priority — Specifies the amount to raise the priority of the destination.
Values 1 to 255

unicast-rt-test

| | |
|--------------------|--|
| Syntax | unicast-rt-test no unicast-rt-test |
| Context | config>filter>redirect-policy>dest |
| Description | <p>This command configures a unicast route test for this destination. A destination is eligible for redirect if a valid unicast route to that destination exists in the routing instance specified by config>filter>redirect-policy>router. The unicast route test is mutually exclusive with other redirect-policy test types.</p> <p>The test cannot be configured if no router is configured for this redirect policy.</p> <p>The no form of the command disables the test.</p> |
| Default | no unicast-rt-test |

url-test

| | |
|--------------------|--|
| Syntax | url-test <i>test-name</i> no url-test <i>test-name</i> |
| Context | config>filter>redirect-policy>dest |
| Description | The context to enable URL test parameters. IP filters can be used to selectively cache some web sites. |
| Default | n/a |
| Parameters | <i>test-name</i> — Specifies the name of the URL test. Allowed values are any string up to 32 characters in length composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, and so on), the entire string must be enclosed within double quotes. |

return-code

| | |
|--------------------|--|
| Syntax | return-code <i>return-code-1</i> [<i>return-code-2</i>] [disable lower-priority <i>priority</i> raise-priority <i>priority</i>] no return-code <i>return-code-1</i> [<i>return-code-2</i>] |
| Context | config>filter>redirect-policy>dest>url-test |
| Description | Return codes are returned when the URL test is performed. Values for the specified range are the return codes which can be given back to the system as a result of the test been performed. |

For example, error code 401 for HTTP is “page not found.” If, while performing this test, the URL is not reachable, you can lower the priority by 10 points so that other means of reaching this destination are prioritized higher than the older one.

| | |
|-------------------|---|
| Default | n/a |
| Parameters | <p><i>return-code-1</i>, <i>return-code-2</i> — Specifies a range of return codes. When the URL test return-code falls within the specified range, the corresponding action is performed.</p> <p>Values <i>return-code-1</i>: 1 to 4294967294 <i>return-code-2</i>: 2 to 4294967295</p> <p>disable — Specifies that the destination may not be used for the amount of time specified in the hold-time command when the return code falls within the specified range.</p> <p>lower-priority <i>priority</i> — Specifies the amount to lower the priority of the destination when the return code falls within the specified range.</p> <p>raise-priority <i>priority</i> — Specifies the amount to raise the priority of the destination when the return code falls within the specified range.</p> |

url

| | |
|--------------------|--|
| Syntax | <p>url <i>url-string</i> [http-version <i>version-string</i>]</p> <p>no url</p> |
| Context | config>filter>redirect-policy>dest>url-test |
| Description | This command specifies the URL to be probed by the URL test. |
| Default | n/a |
| Parameters | <p><i>url-string</i> — Specifies a URL up to 255 characters in length.</p> <p><i>version-string</i> — Specifies the HTTP version, 80 characters in length.</p> |

router

| | |
|--------------------|--|
| Syntax | <p>router <i>router-instance</i></p> <p>router service-name <i>service-name</i></p> <p>no router</p> |
| Context | config>filter>redirect-policy |
| Description | This command enhances VRF support in redirect policies. When a router instance is specified, the configured destination tests are run in the specified router instance, and the PBR action is executed in the specified router instance. If no destination is active or if the hardware does not support PBR action “next-hop router”, action forward will be executed (i.e. routing will be performed in the context of the incoming interface routing instance). |

The **no** form of the command preserves backward-compatibility. Tests always run in the “Base” routing instance context, and the PBR action executes in the routing context of the ingress interface that the filter using this redirect policy is deployed on.

| | |
|-------------------|---|
| Default | no router |
| Parameters | <i>router-instance</i> — Specifies a router instance in the form of router-name or service-id . |
| Values | router-name — Base service-id — Specifies an existing Layer 3 service [1 to 2147483647] <i>service-name</i> — Specifies the name of a configured Layer 3 service. |

shutdown

| | |
|--------------------|--|
| Syntax | [no] shutdown |
| Context | config>filter>redirect-policy config>filter>redirect-policy>destination |
| Description | <p>Administratively enables/disabled (AdminUp/AdminDown) an entity. Downing an entity does not change, reset or remove any configuration settings or statistics. Many objects must be shutdown before they may be deleted.</p> <p>The shutdown command administratively downs an entity. Administratively downing an entity changes the operational state of the entity to down.</p> <p>Unlike other commands and parameters where the default state will not be indicated in the configuration file, shutdown and no shutdown are always indicated in system generated configuration files.</p> <p>The no form of the command puts an entity into the administratively enabled state.</p> |
| Default | no shutdown |

4.5 Show, Clear, Monitor, and Debug Command Reference

- [Command Hierarchies](#)
- [Command Descriptions](#)

4.5.1 Command Hierarchies

- [Show Commands](#)
- [Clear Commands](#)
- [Monitor Commands](#)
- [Debug Commands](#)
- [Tools Commands](#)

4.5.1.1 Show Commands

```
show
  — filter
    — dhcp [filter-id]
    — dhcp6 [filter-id]
    — ip [filter-type filter-type]
    — ip embedded [inactive]
    — ip ip-filter-id embedded [inactive]
    — ip ip-filter-id [detail]
    — ip ip-filter-id associations
    — ip ip-filter-id type entry-type
    — ip ip-filter-id counters [type entry-type] [detail]
    — ip ip-filter-id entry entry-id [counters] [detail]
    — ip ip-filter-id [entry entry-id] effective-action [router | service {service-id | service-name}] [ingress | egress]
    — ipv6 [filter-type filter-type]
    — ipv6 embedded [inactive]
    — ipv6 ipv6-filter-id embedded [inactive]
    — ipv6 ipv6-filter-id [detail]
    — ipv6 ipv6-filter-id associations
    — ipv6 ipv6-filter-id type entry-type
    — ipv6 ipv6-filter-id counters [type entry-type] [detail]
    — ipv6 ipv6-filter-id entry entry-id [counters] [detail]
    — ipv6 ipv6-filter-id [entry entry-id] effective-action [router | service {service-id | service-name}] [ingress | egress]
    — log [bindings]
    — log log-id [match string]
```

- **mac** *mac-filter-id*
- **mac** *mac-filter-id* **associations**
- **mac** *mac-filter-id* [**type** *entry-type*] **counters** [**detail**]
- **mac** [*mac-filter-id*] **embedded** [**inactive**]
- **mac** *mac-filter-id* **entry** *entry-id* [**counters**] [**detail**]
- **mac** [**filter-type** *filter-type*]
- **mac** *mac-filter-id* **type** *entry-type*
- **mac** *mac-filter-id* [**entry** *entry-id*] **effective-action** [**router** | **service** {*service-id* | *service-name*}] [**ingress** | **egress**]
- **match-list**
 - **ip-prefix-list** [*prefix-list-name*]
 - **ip-prefix-list** *prefix-list-name* **references**
 - **ipv6-prefix-list** [*prefix-list-name*]
 - **ipv6-prefix-list** *prefix-list-name* **references**
 - **port-list** [*port-list-name*]
 - **port-list** *port-list-name* **references**
- **redirect-policy** [*redirect-policy-name* {**dest** *ip-address* | **associations**}]
- **system-filter** [**chained-to**]

4.5.1.2 Clear Commands

- ```
clear
 — filter
 — ip filter-id [entry entry-id] [ingress | egress]
```
- ```
    — ipv6 ipv6-filter-id [entry entry-id] [ingress | egress]
```
- ```
 — log log-id
```
- ```
    — mac mac-filter-id [entry entry-id] [ingress | egress]
```

4.5.1.3 Monitor Commands

- ```
monitor
 — filter
 — ip filter-id entry entry-id [interval seconds] [repeat repeat] [absolute | rate]
```
- ```
    — ipv6 ipv6-filter-id entry entry-id [interval seconds] [repeat repeat] [absolute | rate]
```
- ```
 — mac mac-filter-id entry entry-id [interval seconds] [repeat repeat] [absolute | rate]
```

#### 4.5.1.4 Debug Commands

- ```
tools
  — dump
    — filter
      — resources
        — cpm
```
- ```
 — dest-tracking {sap | sdp | ip | ipv6} [detail]
```
- ```
        — egress-pbr [detail]
```


- **http-redirect** [detail]
- **iom** [slot-number]
- **ip** filter-id
- **ipv6** filter-id
- **mac** filter-id
- **sticky-dest**

4.5.1.5 Tools Commands

- ```
tools
 — perform
 — filter
 — ip-filter
 — entry
 — activate-primary-action
 — ipv6-filter
 — entry
 — activate-primary-action
 — mac-filter
 — entry
 — activate-primary-action
 — redirect-policy
 — activate-best-dest
```

## 4.5.2 Command Descriptions

- [Show Commands](#)
- [Clear Commands](#)
- [Monitor Commands](#)
- [Debug Commands](#)

### 4.5.2.1 Show Commands

The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

dhcp

**Syntax**    **dhcp** [filter-id]

---

|                    |                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | show>filter                                                                                                                     |
| <b>Description</b> | This command displays DHCP filter information.                                                                                  |
| <b>Parameters</b>  | <i>filter-id</i> — displays detailed information for the specified filter ID and its filter entries<br><b>Values</b> 1 to 65535 |
| <b>Output</b>      | The following is a sample command output for the command when no filter ID is specified.                                        |

**Sample Output**

```
*B:TechPubs>config# show filter dhcp
=====
DHCP Filters
=====
Filter-Id Applied Description

10 No test-dhcp-filter

Num filter entries: 1
=====
*B:TechPubs>config#

*B:TechPubs>config# show filter dhcp 10
=====
DHCP Filter
=====
Filter-Id : 10 Applied : No
Entries : 0
Description : test-dhcp-filter

Filter Match Criteria

No Match Criteria Found
=====
*B:TechPubs>config#
```

**dhcp6**

|                    |                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>dhcp6</b> [ <i>filter-id</i> ]                                                                                               |
| <b>Context</b>     | show>filter                                                                                                                     |
| <b>Description</b> | This command displays DHCP6 filter information.                                                                                 |
| <b>Parameters</b>  | <i>filter-id</i> — displays detailed information for the specified filter ID and its filter entries<br><b>Values</b> 1 to 65535 |

## ip

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ip</b> [ <b>filter-type</b> <i>filter-type</i> ]<br><b>ip embedded</b> [ <b>inactive</b> ]<br><b>ip</b> <i>ip-filter-id</i> <b>embedded</b> [ <b>inactive</b> ]<br><b>ip</b> <i>ip-filter-id</i> [ <b>detail</b> ]<br><b>ip</b> <i>ip-filter-id</i> <b>associations</b><br><b>ip</b> <i>ip-filter-id</i> <b>type</b> <i>entry-type</i><br><b>ip</b> <i>ip-filter-id</i> <b>counters</b> [ <b>type</b> <i>entry-type</i> ] [ <b>detail</b> ]<br><b>ip</b> <i>ip-filter-id</i> <b>entry</b> <i>entry-id</i> [ <b>counters</b> ] [ <b>detail</b> ]<br><b>ip</b> <i>ip-filter-id</i> [ <b>entry</b> <i>entry-id</i> ] <b>effective-action</b> [ <b>router</b>   <b>service</b> { <i>service-id</i>   <i>service-name</i> }] [ <b>ingress</b>   <b>egress</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Context</b>     | show>filter                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b> | <p>This command displays IPv4 filter information.</p> <p>When <b>effective-action</b> is specified, this command displays what effectively happens to a packet that matches the criteria associated with the IPv4 filter policy.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Parameters</b>  | <p><b>filter-type</b> <i>filter-type</i> — Specifies the type of filter to display.</p> <p><b>Values</b> config, flowspec, host-common, openflow, vsd</p> <p><i>ip-filter-id</i> — Specifies the IPv4 filter policy for which to display information. Values can be expressed in different formats; the following shows decimal integer format.</p> <p><b>Values</b> 1 to 65535</p> <p><b>entry</b> <i>entry-id</i> — Specifies the filter policy entry (of the specified filter policy) for which to display information.</p> <p><b>Values</b> 1 to 65535</p> <p><b>associations</b> — Appends, to the detailed filter policy output, information about where the specified filter policy is applied.</p> <p><b>counters</b> — Displays counter information for the specified filter ID. Egress counters count the packets without Layer 2 encapsulation. Ingress counters count the packets with Layer 2 encapsulation.</p> <p><b>type</b> <i>entry-type</i> — Specifies the type of filter entry to display.</p> <p><b>Values</b> fixed, radius-insert, credit-control-insert, flow-spec, embedded, radius-shared, pcc-rule (applies only to the 7750 SR)</p> <p><b>embedded</b> [<b>inactive</b>] — Displays all, or optionally inactive embeddings. If <i>ip-filter-id</i> is specified, displays embeddings for the IP filter ID.</p> <p><b>effective-action</b> — Displays the action that the system will effectively apply to the packet.</p> <p><b>router</b> — Filters the output and only displays the information for that specific service ("Base" instance).</p> |

**service *service-id*** — Filters the output and only displays the information for the specified service. The specified value must correspond to an existing service in which the filter has been applied.

**service *service-name*** — Filters the output and only displays the information for the specified service. The specified value must correspond to an existing service in which the filter has been applied.

**ingress** — Filters the output and only displays the information for filter policies applied on ingress.

**egress** — Filters the output and only displays the information for filter policies applied on egress.

**Output**     **Show Filter (no filter-id specified)** — The following is a sample output of IPv4 filter information when no filter ID is specified. [Table 53](#) describes the command output fields.

**Sample Output**

```
A:ALA-49# show filter ip
=====
Configured IP Filters Total: 2
=====
Filter-Id Scope Applied Description

5 Template Yes
6 Template Yes

=====
Host Common IP Filters Total: 2
=====
Filter-Id Description

5:P4 Auto-created PCC-Rule Ingress Filter
6:P5 Auto-created PCC-Rule Egress Filter
=====
Num IP filters: 4
=====
```

**Table 53     Filter IP Output Fields (No Filter ID Specified)**

| Label     | Description                                          |
|-----------|------------------------------------------------------|
| Filter Id | The IP filter ID.                                    |
| Scope     | Template<br>The filter policy is of type template.   |
|           | Exclusive<br>The filter policy is of type exclusive. |

**Table 53 Filter IP Output Fields (No Filter ID Specified) (Continued)**

| Label       | Description                                      |
|-------------|--------------------------------------------------|
| Applied     | No<br>The filter policy ID has not been applied. |
|             | Yes<br>The filter policy ID is applied.          |
| Description | The IP filter policy description.                |

**Show Filter (no filter-id specified, embedded keyword specified)** — The following is a sample output of IPv4 filter information when no filter ID is specified but the embedded keyword is specified. [Table 54](#) describes the command output fields.

### Sample Output

```
*A:Dut-C>config>filter# show filter ip embedded
=====
IP Filter embedding
=====
In From Priority Inserted Status

10 2 50 1/1 OK
 1 100 1/2 OK- 1 entry overwritten
20 2 100 0/5 Failed - out of resources
=====
```

**Table 54 Filter IP Output Fields (No Filter ID Specified, Embedded Keyword Specified)**

| Label    | Description                                                                                                                                                                                                                                                       |
|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| In       | Shows embedding filter index.                                                                                                                                                                                                                                     |
| From     | Shows embedded filters included.                                                                                                                                                                                                                                  |
| Priority | Shows priority of embedded filter.                                                                                                                                                                                                                                |
| Inserted | Shows embedded/total number of entries from embedded filter.<br>Status:<br><b>OK</b> —Embedding operation successful, if any entries are overwritten this will also be indicated.<br><b>Failed</b> —Embedding failed, the reason is displayed (out of resources). |

**Show Filter (with filter-id specified)** — The following is a sample output of IPv4 filter information with the filter ID specified. [Table 55](#) describes the command output fields.

### Sample Output

```
*A:dut-a_a>config>filter>ip-filter>entry>action$ show filter ip 2
=====
IP Filter
=====
Filter Id : 2 Applied : No
Scope : Template Def. Action : Drop
System filter : Unchained
Radius Ins Pt : n/a
CrCtl. Ins Pt : n/a
RadSh. Ins Pt : n/a
PccRl. Ins Pt : n/a
Entries : 1
Description : (Not Specified)

Filter Match Criteria : IP

Entry : 1
Description : (Not Specified)
Log Id : n/a
Src. IP : 0.0.0.0/0
Src. Port : n/a
Dest. IP : 0.0.0.0/0
Dest. Port : n/a
Protocol : Undefined Dscp : Undefined
ICMP Type : Undefined ICMP Code : Undefined
Fragment : Off Src Route Opt : Off
Sampling : Off Int. Sampling : On
IP-Option : 0/0 Multiple Option : Off
TCP-syn : Off TCP-ack : Off
Option-pres : Off
Egress PBR : Disabled
Primary Action : Forward (Next Hop VRF)
 Next Hop : 1.2.3.4
 Router : Base
 PBR Target Status : Down
 Extended Action : Remark DSCP "be"
Secondary Action : Forward (Next Hop VRF)
 Next Hop : 3.4.5.6
 Router : 32
 PBR Target Status : Down
 Extended Action : Remark DSCP "ef"
PBR Down Action : Drop (entry-default)
Downloaded Action : None
Dest. Stickiness : None Hold Remain : 0
Ing. Matches : 0 pkts
Egr. Matches : 0 pkts
=====
```

**Table 55 Show Filter IP (with Filter ID Specified) Output Fields**

| Label                 | Description                                                                                                                                         |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| Filter Id             | The IPv4 filter policy ID.                                                                                                                          |
| Applied               | No<br>The filter policy ID has not been applied.                                                                                                    |
|                       | Yes<br>The filter policy ID is applied.                                                                                                             |
| Scope                 | Template<br>The filter policy is of type template.                                                                                                  |
|                       | Exclusive<br>The filter policy is of type exclusive.                                                                                                |
| Def. Action           | Forward<br>The default action for the filter ID for packets that do not match the filter entries is to forward.                                     |
|                       | Drop<br>The default action for the filter ID for packets that do not match the filter entries is to drop.                                           |
| System filter         | Indicates if the filter has been chained to a system filter.                                                                                        |
| Radius Ins Pt         | Indicates the RADIUS insertion point, if any.                                                                                                       |
| CrCtl. Ins Pt         | Indicates the Credit Control insertion point, if any.                                                                                               |
| RadSh. Ins Pt         | Indicates the RADIUS shared insertion point, if any.                                                                                                |
| PccRl. Ins Pt         | Indicates the PCC rule insertion point, if any.                                                                                                     |
| Entries               | The number of entries configured in this filter ID.                                                                                                 |
| Description           | The IPv4 filter policy entry description string.                                                                                                    |
| Filter Match Criteria | IP<br>Indicates the filter is an IPv4 filter policy.                                                                                                |
| Entry                 | The filter entry ID. If the filter entry ID indicates the entry is (Inactive), then the filter entry is incomplete as no action has been specified. |
| Description           | The IPv4 filter policy entry description string.                                                                                                    |
| Log Id                | The filter log ID.                                                                                                                                  |
| Src. IP               | The source IPv4 address and prefix length match criterion. "0.0.0.0/0" indicates no criterion specified for the filter entry.                       |

**Table 55 Show Filter IP (with Filter ID Specified) Output Fields (Continued)**

| Label         | Description                                                                                                                          |
|---------------|--------------------------------------------------------------------------------------------------------------------------------------|
| Src. Port     | The source TCP, UDP, or SCTP port number, port range, or port match list.                                                            |
| Dest. IP      | The destination IPv4 address and prefix length match criterion. "0.0.0.0/0" indicates no criterion specified for the filter entry.   |
| Dest. Port    | The destination TCP, UDP, or SCTP port number, port range, or port match list.                                                       |
| Protocol      | The protocol for the match criteria. Undefined indicates no protocol specified.                                                      |
| Dscp          | The DiffServ Code Point (DSCP) name.                                                                                                 |
| ICMP Type     | The ICMP type match criterion. Undefined indicates no ICMP type specified.                                                           |
| ICMP Code     | The ICMP code field in the ICMP header of an IPv4 packet.                                                                            |
| Fragment      | False<br>Configures a match on all non-fragmented IPv4 packets.                                                                      |
|               | True<br>Configures a match on all fragmented IPv4 packets.                                                                           |
|               | Off<br>Fragments are not a matching criteria. All fragments and non-fragments implicitly match.                                      |
| Src Route Opt | Indicates if the source route option has been set.                                                                                   |
| Sampling      | Off<br>Specifies that traffic sampling is disabled.                                                                                  |
|               | On<br>Specifies that traffic matching the associated IPv4 filter entry is sampled.                                                   |
| Int. Sampling | Off<br>Interface traffic sampling is disabled.                                                                                       |
|               | On<br>Interface traffic sampling is enabled.                                                                                         |
| IP-Option     | Specifies matching packets with a specific IPv4 option or a range of IPv4 options in the IPv4 header for IPv4 filter match criteria. |



**Table 55 Show Filter IP (with Filter ID Specified) Output Fields (Continued)**

| Label             | Description                                                                                                                              |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| Multiple Option   | Off<br>The option fields are not checked.                                                                                                |
|                   | On<br>Packets containing one or more option fields in the IPv4 header will be used as IPv4 filter match criteria.                        |
| TCP-syn           | False<br>Configures a match on packets with the SYN flag set to false.                                                                   |
|                   | True<br>Configured a match on packets with the SYN flag set to true.                                                                     |
|                   | Off<br>The state of the TCP SYN flag is not considered as part of the match criteria.                                                    |
| TCP-ack           | False<br>Configures a match on packets with the ACK flag set to false.                                                                   |
|                   | True<br>Configures a match on packets with the ACK flag set to true.                                                                     |
|                   | Off<br>The state of the TCP ACK flag is not considered as part of the match criteria.                                                    |
| Option-present    | Off<br>Specifies not to search for packets that contain the option field or have an option field of zero.                                |
|                   | On<br>Matches packets that contain the option field or have an option field of zero be used as IPv4 filter match criteria.               |
| Egress PBR        | Indicates if the <b>egress-pbr</b> flag is set for this entry.                                                                           |
| Primary Action    | Indicates the configured action, if any. Indented sub-labels in the show output provide configured parameters for this action            |
| Secondary Action  | Indicates the configured secondary action, if any. Indented sub-labels in the show output provide configured parameters for this action. |
| PBR Target Status | The status of the target of the primary or secondary action based on simple checks.                                                      |

**Table 55** Show Filter IP (with Filter ID Specified) Output Fields (Continued)

| Label             | Description                                                                                                                                                                       |
|-------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Extended Action   | Indicates the configured extended action, if any.                                                                                                                                 |
| PBR Down Action   | Indicates the action to take when the target is down. Packets that match the entry criteria will be subject to the PBR Down Action in case the target of the main action is down. |
| Downloaded Action | The action downloaded by CPM to IOM.                                                                                                                                              |
| Dest. Stickiness  | Indicates whether stickiness is configured.                                                                                                                                       |
| Hold Remain       | The stickiness timer.                                                                                                                                                             |
| Ing. Matches      | The number of ingress filter matches/hits for the filter entry.                                                                                                                   |
| Egr. Matches      | The number of egress filter matches/hits for the filter entry.                                                                                                                    |

**Show Filter Associations** — The following is a sample output of IPv4 filter information when the **associations** keyword is specified. [Table 56](#) describes the command output fields.

**Sample Output**

```

A:ALA-49# show filter ip 1 associations
=====
IP Filter
=====
Filter Id : 4 Applied : Yes
Scope : Template Def. Action : Drop
System filter : Unchained
Radius Ins Pt : n/a
CrCtl. Ins Pt : n/a
RadSh. Ins Pt : n/a
PccRl. Ins Pt : n/a
Entries : 1
Description : (Not Specified)

Filter Association : IP

Service Id : 2 Type : VPLS
- SAP 1/2/2 (Ingress)

Filter associated with IOM: 1
=====

```

**Table 56** Filter IP Associations Output Fields

| Label     | Description                |
|-----------|----------------------------|
| Filter Id | The IPv4 filter policy ID. |

**Table 56 Filter IP Associations Output Fields (Continued)**

| Label              | Description                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| Applied            | No<br>The filter policy ID has not been applied.                                                                                              |
|                    | Yes<br>The filter policy ID is applied.                                                                                                       |
| Scope              | Template<br>The filter policy is of type Template.                                                                                            |
|                    | Exclusive<br>The filter policy is of type Exclusive.                                                                                          |
| Def. Action        | Forward<br>The default action for the filter ID for packets that do not match the filter entries is to forward.                               |
|                    | Drop<br>The default action for the filter ID for packets that do not match the filter entries is to drop.                                     |
| System filter      | Indicates if the filter has been chained to a system filter.                                                                                  |
| Radius Ins Pt      | Indicates the RADIUS insertion point, if any.                                                                                                 |
| CrCtl. Ins Pt      | Indicates the Credit Control insertion point, if any.                                                                                         |
| RadSh. Ins Pt      | Indicates the RADIUS shared insertion point, if any.                                                                                          |
| PccRI. Ins Pt      | Indicates the PCC rule insertion point, if any.                                                                                               |
| Entries            | The number of entries configured in this filter ID.                                                                                           |
| Description        | The IPv4 filter policy description.                                                                                                           |
| Filter Association | Indicates the filter is an IPv4 filter policy.                                                                                                |
| Service Id         | The service ID on which the filter policy ID is applied. The output also provides a list of service points where the filter has been applied. |
| Type               | The type of service of the service ID.                                                                                                        |
| (Ingress)          | The filter policy ID is applied as an ingress filter policy on the interface.                                                                 |
| (Egress)           | The filter policy ID is applied as an egress filter policy on the interface.                                                                  |

**Show Filter Counters** — The following is a sample output of IPv4 filter information when the **counters** keyword is specified. [Table 57](#) describes the command output fields.

Egress counters count the packets without Layer 2 encapsulation. Ingress counters count the packets with Layer 2 encapsulation.

### Sample Output

```
*A:ALA-48# show filter ip 100 counters
=====
IP Filter
=====
Filter Id : 4 Applied : Yes
Scope : Template Def. Action : Drop
System filter : Unchained
Radius Ins Pt : n/a
CrCtl. Ins Pt : n/a
RadSh. Ins Pt : n/a
PccRl. Ins Pt : n/a
Entries : 1
Description : (Not Specified)

Filter Match Criteria : IP

Entry : 4001
Ing. Matches : 9788619 pkts (978861900 bytes)
Egr. Matches : 9788619 pkts (978861900 bytes)
=====
```

**Table 57** Filter IP Counters Output Field Descriptions

| Label     | Description                                          |
|-----------|------------------------------------------------------|
| Filter Id | The IPv4 filter policy ID.                           |
| Applied   | No<br>The filter policy ID has not been applied.     |
|           | Yes<br>The filter policy ID is applied.              |
| Scope     | Template<br>The filter policy is of type Template.   |
|           | Exclusive<br>The filter policy is of type Exclusive. |

**Table 57 Filter IP Counters Output Field Descriptions (Continued)**

| Label                 | Description                                                                                                                                         |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| Def. Action           | Forward<br>The default action for the filter ID for packets that do not match the filter entries is to forward.                                     |
|                       | Drop<br>The default action for the filter ID for packets that do not match the filter entries is to drop.                                           |
| System filter         | Indicates if the filter has been chained to a system filter.                                                                                        |
| Radius Ins Pt         | Indicates the RADIUS insertion point, if any.                                                                                                       |
| CrCtl. Ins Pt         | Indicates the Credit Control insertion point, if any.                                                                                               |
| RadSh. Ins Pt         | Indicates the RADIUS shared insertion point, if any.                                                                                                |
| PccRI. Ins Pt         | Indicates the PCC rule insertion point, if any.                                                                                                     |
| Entries               | The number of entries configured in this filter ID.                                                                                                 |
| Description           | The IPv4 filter policy description.                                                                                                                 |
| Filter Match Criteria | IP<br>Indicates the filter is an IPv4 filter policy.                                                                                                |
| Entry                 | The filter entry ID. If the filter entry ID indicates the entry is (Inactive), then the filter entry is incomplete as no action has been specified. |
| Ing. Matches          | The number of ingress filter matches/hits for the filter entry.                                                                                     |
| Egr. Matches          | The number of egress filter matches/hits for the filter entry.                                                                                      |

**Show Filter IP Output (with effective-action specified)** — The following is a sample output of IPv4 filter information when the **effective-action** keyword is specified. [Table 58](#) describes the command output fields.

If the main action (either primary or secondary) cannot be performed, a reason will be given. This will be displayed on the same line as the Effective Action. The reason codes as currently defined are:

- action not supported in L2 service
- action not supported in L3 service
- action not supported on egress
- destination not reachable
- egress-pbr is off
- egress-pbr is on
- entry-default

- filter-default-action
- not POS unnumbered interface
- pbr-down-action-override
- target does not exist

### Sample Output

```

show filter ip 1 effective-action
=====
IP Filter
=====
Filter Id : 1 Applied : Yes
Scope : Template Def. Action : Drop
Entries : 1
Description : (Not Specified)

Entry : 1

Stickiness : No
PBR Dwn Act Override: None
PBR Down Action : Drop (entry-default)

Configuration
Primary Action : Forward (SAP)
 Next Hop : 1/1/2
 Service Id : 10
Secondary Action : None

Status
Target status based on extended checks
 Primary Action : Down
 Secondary Action : None
Downloaded Action : Primary
Stickiness Timer : Not Running

Effective Action based on application context
Service Id : 10 Type : VPLS
Ingress
 Effective Action: Drop (entry-default)
=====

```

**Table 58** Show Filter IP effective-action Output Field Descriptions

| Label     | Description                                      |
|-----------|--------------------------------------------------|
| Filter Id | The IPv4 filter policy ID.                       |
| Applied   | No<br>The filter policy ID has not been applied. |
|           | Yes<br>The filter policy ID is applied.          |

**Table 58 Show Filter IP effective-action Output Field Descriptions**

| Label                                         | Description                                                                                                                                                                       |
|-----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Scope                                         | Template<br>The filter policy is of type Template.                                                                                                                                |
|                                               | Exclusive<br>The filter policy is of type Exclusive.                                                                                                                              |
| Def. Action                                   | Forward<br>The default action for the filter ID for packets that do not match the filter entries is to forward.                                                                   |
|                                               | Drop<br>The default action for the filter ID for packets that do not match the filter entries is to drop.                                                                         |
| Entries                                       | The number of entries configured in this filter ID.                                                                                                                               |
| Description                                   | The IPv4 filter policy description.                                                                                                                                               |
| Entry                                         | The filter entry ID. If the filter entry ID indicates the entry is (Inactive), then the filter entry is incomplete as no action has been specified.                               |
| Stickiness                                    | No<br>Stickiness is not configured.                                                                                                                                               |
|                                               | Yes<br>Stickiness is configured.                                                                                                                                                  |
| PBR Dwn Act Override                          | Indicates whether or not the action to take when the PBR target is down has been overridden.                                                                                      |
| PBR Down Action                               | Indicates the action to take when the target is down. Packets that match the entry criteria will be subject to the PBR Down Action in case the target of the main action is down. |
| Configuration                                 | Section of the output providing information on the configured parameters.                                                                                                         |
| Primary Action                                | The configured action, if any. Indented sub-labels in the show output provide configured parameters for this action.                                                              |
| Secondary Action                              | The configured secondary action, if any. Indented sub-labels in the show output provide configured parameters for this action.                                                    |
| Status/Target status based on extended checks | Section of the output providing information on the operational status of certain parameters.                                                                                      |

**Table 58** Show Filter IP effective-action Output Field Descriptions

| Label                                         | Description                                                                                                                                   |
|-----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| Primary Action                                | The status of the target of the primary action, if any configured, based on extended checks.                                                  |
| Secondary Action                              | The status of the target of the secondary action, if any configured, based on extended checks.                                                |
| Downloaded Action                             | The action downloaded by the CPM to the IOM.                                                                                                  |
| Stickiness Timer                              | The status of the stickiness timer, if any.                                                                                                   |
| Effective Action based on application context | Section of the output providing the effective action, in the context of services, that a packet matching the criteria will be subject to.     |
| Service Id                                    | The service ID on which the filter policy ID is applied. The output also provides a list of service points where the filter has been applied. |
| Type                                          | The service type in which the service has been applied.                                                                                       |
| Ingress/Egress                                | The direction in which the service has been applied.                                                                                          |
| Effective Action                              | The effective action that the packet will be subject to.                                                                                      |
| Extended Action                               | Indicates the configured extended action, if any.                                                                                             |

## ipv6

**Syntax**

```

ipv6 [filter-type filter-type]
ipv6 embedded [inactive]
ipv6 ipv6-filter-id embedded [inactive]
ipv6 ipv6-filter-id [detail]
ipv6 ipv6-filter-id associations
ipv6 ipv6-filter-id type entry-type
ipv6 ipv6-filter-id counters [type entry-type] [detail]
ipv6 ipv6-filter-id entry entry-id [counters] [detail]
ipv6 ipv6-filter-id [entry entry-id] effective-action [router | service {service-id | service-name}] [ingress | egress]

```

**Context** show>filter

**Description** This command shows IPv6 filter information.

When **effective-action** is specified, this command displays what effectively happens to a packet that matches the criteria associated with the IPv6 filter policy.



- Parameters**
- ipv6-filter-id*** — specifies the IPv6 filter policy for which to display information. Values can be expressed in different formats. The following only shows decimal integer format.  
**Values** 1 to 65535
  - entry entry-id*** — specifies the filter policy entry (of the specified filter policy) for which only to display information  
**Values** 1 to 65535
  - filter-type filter-type*** — specifies the type of filter to display  
**Values** config, flowspec, host-common, openflow, vsd.
  - associations*** — appends, to the detailed filter policy output, information as to where the specified filter policy ID is applied
  - counters*** — displays counter information. Egress counters count the packets without Layer 2 encapsulation. Ingress counters count the packets with Layer 2 encapsulation.
  - embedded [inactive]*** — shows all embeddings, optionally shows inactive embedding only, if *ipv6-filter-id* is not specified shows all embedded filters
  - type entry-type*** — specifies type of filter entry to display:  
**Values** fixed, radius-insert, credit-control-insert, embedded, radius-shared
  - effective-action*** — Displays the action that the system will effectively apply to the packet.
  - router*** — Filters the output and only displays the information for that specific service ("Base" instance).
  - service service-id*** — Filters the output and only displays the information for the specified service. The specified value should correspond to an existing service in which the filter has been applied.
  - service service-name*** — Filters the output and only displays the information for the specified service. The specified value should correspond to an existing service in which the filter has been applied.
  - ingress*** — Filters the output and only displays the information for filter policies applied on ingress.
  - egress*** — Filters the output and only displays the information for filter policies applied on egress.

**Output** **Show Filter (no filter-id specified)** — The following output is an example of IPv6 filter information when no filter ID is specified, and [Table 59](#) describes the fields.

### Sample Output

```
A:ALA-48# show filter ipv6
=====
IP Filters
=====
Filter-Id Scope Applied Description
```

```

100 Template Yes test
200 Exclusive Yes

Num IPv6 filters: 2
=====
A:ALA-48# show filter ipv6 embedded
=====
IP Filter embedding
=====
In From Priority Inserted Status

10 2 50 1/1 OK
 1 100 1/2 OK- 1 entry overwritten

20 2 100 0/5 Failed - out of resources
=====
A:ALA-48#
=====
Configured IP Filters Total: 4
=====
Filter-Id Scope Applied Description

1 Template No
5 Exclusive No
10 Template Yes
100 Embedded N/A
=====
System IP Filters Total: 1
=====
Filter-Id Description

_tmnx_ofs_test of-switch 'test' embedded filter

Num IP filters: 5
=====

```

**Table 59** Filter IPv6 Output Fields

| Label     | Description                                          |
|-----------|------------------------------------------------------|
| Filter Id | The IPv6 filter ID.                                  |
| Scope     | Template<br>The filter policy is of type template.   |
|           | Exclusive<br>The filter policy is of type exclusive. |
| Applied   | No<br>The filter policy ID has not been applied.     |
|           | Yes<br>The filter policy ID is applied.              |

**Table 59 Filter IPv6 Output Fields (Continued)**

| Label       | Description                                                                                                                                                                                                                                                       |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Description | The IPv6 filter policy description.                                                                                                                                                                                                                               |
| In          | Shows embedding filter index.                                                                                                                                                                                                                                     |
| From        | Shows embedded filters included.                                                                                                                                                                                                                                  |
| Priority    | Shows priority of embedded filter.                                                                                                                                                                                                                                |
| Inserted    | Shows embedded/total number of entries from embedded filter.<br>Status:<br><b>OK</b> —Embedding operation successful, if any entries are overwritten this will also be indicated.<br><b>Failed</b> —Embedding failed, the reason is displayed (out of resources). |

**Show Filter (with filter-id specified)** — The following output is an example of IPv6 filter information when filter-id is specified, and [Table 60](#) describes the fields.

### Sample Output

```
A:ALA-48# show filter ipv6 100
=====
IPv6 Filter
=====
Filter Id : 100 Applied : Yes
Scope : Template Def. Action : Forward
Entries : 1
Description : test

Filter Match Criteria : IPv6

Entry : 10
Log Id : 101
Src. IP : ::/0 Src. Port : None
Dest. IP : ::/0 Dest. Port : None
Next Header : Undefined Dscp : Undefined
ICMP Type : Undefined ICMP Code : Undefined
TCP-syn : Off TCP-ack : Off
Match action : Drop
Ing. Matches : 0 Egr. Matches : 0
=====
A:ALA-48#
```

**Table 60 Filter IPv6 with Filter-ID Specified Output Fields**

| Label     | Description                |
|-----------|----------------------------|
| Filter Id | The IPv6 filter policy ID. |

**Table 60 Filter IPv6 with Filter-ID Specified Output Fields (Continued)**

| Label                 | Description                                                                                                                                         |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| Scope                 | Template<br>The filter policy is of type template.                                                                                                  |
|                       | Exclusive<br>The filter policy is of type exclusive.                                                                                                |
| Entries               | The number of entries configured in this filter ID.                                                                                                 |
| Description           | The IPv6 filter policy description.                                                                                                                 |
| Applied               | No<br>The filter policy ID has not been applied.                                                                                                    |
|                       | Yes<br>The filter policy ID is applied.                                                                                                             |
| Def. Action           | Forward<br>The default action for the filter ID for packets that do not match the filter entries is to forward.                                     |
|                       | Drop<br>The default action for the filter ID for packets that do not match the filter entries is to drop.                                           |
| Filter Match Criteria | IP<br>Indicates the filter is an IPv6 filter policy.                                                                                                |
| Entry                 | The filter entry ID. If the filter entry ID indicates the entry is (Inactive), then the filter entry is incomplete as no action has been specified. |
| Log Id                | The filter log ID.                                                                                                                                  |
| Src. IP               | The source IPv6 address and mask match criterion.<br>“::/0” indicates no criterion specified for the filter entry.                                  |
| Dest. IP              | The destination IPv6 address and mask match criterion.<br>“::/0” indicates no criterion specified for the filter entry.                             |
| Protocol              | The protocol ID for the match criteria. Undefined indicates no protocol specified.                                                                  |
| ICMP Type             | The ICMP type match criterion. Undefined indicates no ICMP type specified.                                                                          |

**Table 60 Filter IPv6 with Filter-ID Specified Output Fields (Continued)**

| Label        | Description                                                                                                                                                                                                                                                                          |
|--------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Fragment     | False<br>Configures a match on all non-fragmented IPv6 packets.                                                                                                                                                                                                                      |
|              | True<br>Configures a match on all fragmented IPv6 packets.                                                                                                                                                                                                                           |
|              | Off<br>Fragments are not a matching criteria. All fragments and nonfragments implicitly match.                                                                                                                                                                                       |
| Sampling     | Off<br>Specifies that traffic sampling is disabled.                                                                                                                                                                                                                                  |
|              | On<br>Specifies that traffic matching the associated IPv6 filter entry is sampled.                                                                                                                                                                                                   |
| IP-Option    | Specifies matching packets with a specific IPv6 option or a range of IPv6 options in the IPv6 header for IPv6 filter match criteria.                                                                                                                                                 |
| TCP-syn      | False<br>Configures a match on packets with the SYN flag set to false.                                                                                                                                                                                                               |
|              | True<br>Configured a match on packets with the SYN flag set to true.                                                                                                                                                                                                                 |
|              | Off<br>The state of the TCP SYN flag is not considered as part of the match criteria                                                                                                                                                                                                 |
| Match action | Default<br>The filter does not have an explicit forward or drop match action specified. If the filter entry ID indicates the entry is (Inactive), then the filter entry is incomplete as no action has been specified.                                                               |
|              | Drop<br>Drop packets matching the filter entry.                                                                                                                                                                                                                                      |
|              | Forward<br>The explicit action to perform is forwarding of the packet. If the action is Forward, then if configured, the next hop information should be displayed, including nexthop: <i>IPv6 address</i> , Indirect: <i>IPv6 address</i> or Interface: <i>IPv6 interface name</i> . |
| Ing. Matches | The number of ingress filter matches/hits for the filter entry.                                                                                                                                                                                                                      |

**Table 60 Filter IPv6 with Filter-ID Specified Output Fields (Continued)**

| Label             | Description                                                                                                                |
|-------------------|----------------------------------------------------------------------------------------------------------------------------|
| Src. Port         | The source TCP, UDP, or SCTP port number, port range, or port match list.                                                  |
| Dest. Port        | The destination TCP, UDP, or SCTP port number, port range, or port match list.                                             |
| Dscp              | The DiffServ Code Point (DSCP) name.                                                                                       |
| ICMP Code         | The ICMP code field in the ICMP header of an IPv6 packet.                                                                  |
| Option-present    | Off<br>Specifies not to search for packets that contain the option field or have an option field of zero.                  |
|                   | On<br>Matches packets that contain the option field or have an option field of zero be used as IPv6 filter match criteria. |
| Int. Sampling     | Off<br>Interface traffic sampling is disabled.                                                                             |
|                   | On<br>Interface traffic sampling is enabled.                                                                               |
| Multiple Option   | Off<br>The option fields are not checked.                                                                                  |
|                   | On<br>Packets containing one or more option fields in the IPv6 header will be used as IPv6 filter match criteria.          |
| TCP-ack           | False<br>Configures a match on packets with the ACK flag set to false.                                                     |
|                   | True<br>Configured a match on packets with the ACK flag set to true.                                                       |
|                   | Off<br>The state of the TCP ACK flag is not considered as part of the match criteria.                                      |
| Egr. Matches      | The number of egress filter matches/hits for the filter entry.                                                             |
| Ing. Rate-limiter | The number of offered, forwarded, and dropped packet matches for the filter entry.                                         |

**Show Filter Associations** — The following output is an example of IPv6 filter information when the **associations** keyword is specified, and [Table 61](#) describes the fields.

### Sample Output

```
A:ALA-48# show filter ipv6 1 associations
=====
IPv6 Filter
=====
Filter Id : 1 Applied : Yes
Scope : Template Def. Action : Drop
Entries : 1

Filter Association : IPv6

Service Id : 2000 Type : IES
- SAP 1/1/1:2000 (Ingress)
=====
Filter Match Criteria : IPv6

Entry : 10
Log Id : 101
Src. IP : ::/0 Src. Port : None
Dest. IP : ::/0 Dest. Port : None
Next Header : Undefined Dscp : Undefined
ICMP Type : Undefined ICMP Code : Undefined
TCP-syn : Off TCP-ack : Off
Match action : Drop
Ing. Matches : 0 Egr. Matches : 0
=====
A:ALA-48#
```

**Table 61** Filter IPv6 Associations Output Fields

| Label     | Description                                          |
|-----------|------------------------------------------------------|
| Filter Id | The IPv6 filter policy ID.                           |
| Scope     | Template<br>The filter policy is of type Template.   |
|           | Exclusive<br>The filter policy is of type Exclusive. |
| Entries   | The number of entries configured in this filter ID.  |
| Applied   | No<br>The filter policy ID has not been applied.     |
|           | Yes<br>The filter policy ID is applied.              |

**Table 61 Filter IPv6 Associations Output Fields (Continued)**

| Label       | Description                                                                                                                                   |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| Def. Action | Forward<br>The default action for the filter ID for packets that do not match the filter entries is to forward.                               |
|             | Drop<br>The default action for the filter ID for packets that do not match the filter entries is to drop.                                     |
| Service Id  | The service ID on which the filter policy ID is applied. The output also provides a list of service points where the filter has been applied. |
| SAP         | The Service Access Point on which the filter policy ID is applied.                                                                            |
| (Ingress)   | The filter policy ID is applied as an ingress filter policy on the interface.                                                                 |
| (Egress)    | The filter policy ID is applied as an egress filter policy on the interface.                                                                  |
| Type        | The type of service of the service ID.                                                                                                        |
| Entry       | The filter entry ID. If the filter entry ID indicates the entry is Inactive, the filter entry is incomplete, no action was specified.         |
| Log Id      | The filter log ID.                                                                                                                            |
| Src. IP     | The source IPv6 address and mask match criterion.<br>"0.0.0.0/0" indicates no criterion specified for the filter entry.                       |
| Dest. IP    | The destination IPv6 address and mask match criterion.<br>"0.0.0.0/0" indicates no criterion specified for the filter entry.                  |
| Protocol    | The protocol ID for the match criteria. Undefined indicates no protocol specified.                                                            |
| ICMP Type   | The ICMP type match criterion. Undefined indicates no ICMP type specified.                                                                    |
| Fragment    | False<br>Configures a match on all non-fragmented IPv6 packets.                                                                               |
|             | True<br>Configures a match on all fragmented IPv6 packets.                                                                                    |
|             | Off<br>Fragments are not a matching criteria. All fragments and nonfragments implicitly match.                                                |



**Table 61 Filter IPv6 Associations Output Fields (Continued)**

| Label        | Description                                                                                                                                                                                                                                                       |
|--------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Sampling     | Off<br>Specifies that traffic sampling is disabled.                                                                                                                                                                                                               |
|              | On<br>Specifies that traffic matching the associated IPv6 filter entry is sampled.                                                                                                                                                                                |
| IP-Option    | Specifies matching packets with a specific IPv6 option or a range of IPv6 options in the IPv6 header for IPv6 filter match criteria.                                                                                                                              |
| TCP-syn      | False<br>Configures a match on packets with the SYN flag set to false.                                                                                                                                                                                            |
|              | True<br>Configures a match on packets with the SYN flag set to true.                                                                                                                                                                                              |
|              | Off<br>The state of the TCP SYN flag is not considered as part of the match criteria.                                                                                                                                                                             |
| Match action | Default<br>The filter does not have an explicit forward or drop match action specified. If the filter entry ID indicates the entry is inactive, the filter entry is incomplete, no action was specified.                                                          |
|              | Drop<br>Drop packets matching the filter entry.                                                                                                                                                                                                                   |
|              | Forward<br>The explicit action to perform is forwarding of the packet. If the action is Forward, then if configured the nexthop information should be displayed, including Nexthop: <IPv6 address>, Indirect: <IPv6 address> or Interface: <IPv6 interface name>. |
| Ing. Matches | The number of ingress filter matches/hits for the filter entry.                                                                                                                                                                                                   |
| Src. Port    | The source TCP, UDP, or SCTP port number, port range, or port match list.                                                                                                                                                                                         |
| Dest. Port   | The destination TCP, UDP, or SCTP port number, port range, or port match list.                                                                                                                                                                                    |
| Dscp         | The DiffServ Code Point (DSCP) name.                                                                                                                                                                                                                              |
| ICMP Code    | The ICMP code field in the ICMP header of an IPv6 packet.                                                                                                                                                                                                         |

**Table 61 Filter IPv6 Associations Output Fields (Continued)**

| Label           | Description                                                                                                                |
|-----------------|----------------------------------------------------------------------------------------------------------------------------|
| Option-present  | Off<br>Specifies not to search for packets that contain the option field or have an option field of zero.                  |
|                 | On<br>Matches packets that contain the option field or have an option field of zero be used as IPv6 filter match criteria. |
| Int. Sampling   | Off<br>Interface traffic sampling is disabled.                                                                             |
|                 | On<br>Interface traffic sampling is enabled.                                                                               |
| Multiple Option | Off<br>The option fields are not checked.                                                                                  |
|                 | On<br>Packets containing one or more option fields in the IPv6 header will be used as IPv6 filter match criteria.          |
| TCP-ack         | False<br>Configures a match on packets with the ACK flag set to false.                                                     |
|                 | True<br>Configured a match on packets with the ACK flag set to true.                                                       |
|                 | Off<br>The state of the TCP ACK flag is not considered as part of the match criteria.                                      |
| Egr. Matches    | The number of egress filter matches/hits for the filter entry.                                                             |

**Show Filter Counters** — The following output is an example of IPv6 filter information when the **counters** keyword is specified, and [Table 62](#) describes the output fields.

Egress count the packets without Layer 2 encapsulation. Ingress counters count the packets with Layer 2 encapsulation.

### Sample Output

```
A:ALA-48# show filter ipv6 8 counters
=====
IPv6 Filter
=====
Filter Id : 8 Applied : Yes
```

```

Scope : Template Def. Action : Forward
Entries : 4
Description : Description for Ipv6 Filter Policy id # 8

Filter Match Criteria : IPv6

Entry : 5
Ing. Matches : 0 pkts
Egr. Matches : 0 pkts

Entry : 6
Ing. Matches : 0 pkts
Egr. Matches : 0 pkts

Entry : 8
Ing. Matches : 160 pkts (14400 bytes)
Egr. Matches : 80 pkts (6880 bytes)

Entry : 10
Ing. Matches : 80 pkts (7200 bytes)
Egr. Matches : 80 pkts (6880 bytes)

=====
A:ALA-48#

```

**Table 62** Filter IPv6 Counters Output Fields

| Label                  | Description                                                                                                     |
|------------------------|-----------------------------------------------------------------------------------------------------------------|
| IP Filter<br>Filter Id | The IPv6 filter policy ID.                                                                                      |
| Scope                  | Template<br>The filter policy is of type template.                                                              |
|                        | Exclusive<br>The filter policy is of type exclusive.                                                            |
| Applied                | No<br>The filter policy ID has not been applied.                                                                |
|                        | Yes<br>The filter policy ID is applied.                                                                         |
| Def. Action            | Forward<br>The default action for the filter ID for packets that do not match the filter entries is to forward. |
|                        | Drop<br>The default action for the filter ID for packets that do not match the filter entries is to drop.       |

**Table 62 Filter IPv6 Counters Output Fields (Continued)**

| Label                 | Description                                                                                                                                         |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| Filter Match Criteria | IP<br>Indicates the filter is an IPv6 filter policy.                                                                                                |
| Entry                 | The filter entry ID. If the filter entry ID indicates the entry is (Inactive), then the filter entry is incomplete as no action has been specified. |
| Ing. Matches          | The number of ingress filter matches/hits for the filter entry.                                                                                     |
| Egr. Matches          | The number of egress filter matches/hits for the filter entry.                                                                                      |

Egress count the packets without Layer 2 encapsulation. Ingress counters count the packets with Layer 2 encapsulation.

**Show Filter IPv6 Output (with effective-action specified)** — The following is a sample output of IPv6 filter information when the **effective-action** keyword is specified. [Table 63](#) describes the command output fields.

If the main action (either primary or secondary) cannot be performed, a reason will be given. This will be displayed on the same line as the Effective Action. The reason codes as currently defined are:

- action not supported in L2 service
- action not supported in L3 service
- action not supported on egress
- destination not reachable
- egress-pbr is off
- egress-pbr is on
- entry-default
- filter-default-action
- pbr-down-action-override
- target does not exist

### Sample Output

```
show filter ipv6 10 entry 10 effective-action
=====
IPv6 Filter
=====
Filter Id : 10 Applied : Yes
Scope : Template Def. Action : Drop
Entries : 8
Description : (Not Specified)

Entry : 10
```

```

Origin : Fixed - overwrites embedded filter 30 entry 5
Egress PBR : Disabled
Stickiness : No
PBR Dwn Act Override: None
PBR Down Action : Drop (entry-default)

Configuration
Primary Action : Forward (Next Hop VRF)
 Next Hop : 3ffe:0:a0a:a01:: (Indirect)
 Router : Base
 Extended Action : Remark DSCP "cp51"
Secondary Action : Forward (Next Hop VRF)
 Next Hop : 3ffe:0:1414:1401:: (Indirect)
 Router : Base
 Extended Action : Remark DSCP "cp31"

Status
Target status based on extended checks
 Primary Action : Up
 Secondary Action : Up
Downloaded Action : Primary
Stickiness Timer : Not Running

Effective Action based on application context
Service Id : 100 Type : IES
 Ingress
 Effective Action: Primary
 Extended Action : Performed
Service Id : N/A Type : Base Router
 Egress
 Effective Action: Primary
 Extended Action : Performed
=====

```

**Table 63**      **Show Filter IPv6 effective-action Output Field Descriptions**

| Label     | Description                                          |
|-----------|------------------------------------------------------|
| Filter Id | The IPv6 filter policy ID.                           |
| Applied   | No<br>The filter policy ID has not been applied.     |
|           | Yes<br>The filter policy ID is applied.              |
| Scope     | Template<br>The filter policy is of type Template.   |
|           | Exclusive<br>The filter policy is of type Exclusive. |

**Table 63 Show Filter IPv6 effective-action Output Field Descriptions**

| Label                                         | Description                                                                                                                                                                       |
|-----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Def. Action                                   | Forward<br>The default action for the filter ID for packets that do not match the filter entries is to forward.                                                                   |
|                                               | Drop<br>The default action for the filter ID for packets that do not match the filter entries is to drop.                                                                         |
| Entries                                       | The number of entries configured in this filter ID.                                                                                                                               |
| Description                                   | The IPv6 filter policy description.                                                                                                                                               |
| Entry                                         | The filter entry ID. If the filter entry ID indicates the entry is (Inactive), then the filter entry is incomplete as no action has been specified.                               |
| Origin                                        | The type of filter entry.                                                                                                                                                         |
| Egress PBR                                    | Indicates if the <b>egress-pbr</b> flag is set for this entry.                                                                                                                    |
| Stickiness                                    | No<br>Stickiness is not configured.                                                                                                                                               |
|                                               | Yes<br>Stickiness is configured.                                                                                                                                                  |
| PBR Dwn Act Override                          | Indicates whether or not the action to take when the PBR target is down has been overridden.                                                                                      |
| PBR Down Action                               | Indicates the action to take when the target is down. Packets that match the entry criteria will be subject to the PBR Down Action in case the target of the main action is down. |
| Configuration                                 | Section of the output providing information on the configured parameters.                                                                                                         |
| Primary Action                                | The configured action, if any. Indented sub-labels in the show output provide configured parameters for this action.                                                              |
| Secondary Action                              | The configured secondary action, if any. Indented sub-labels in the show output provide configured parameters for this action.                                                    |
| Status/Target status based on extended checks | Section of the output providing information on the operational status of certain parameters.                                                                                      |
| Primary Action                                | The status of the target of the primary action, if configured, based on extended checks.                                                                                          |

**Table 63 Show Filter IPv6 effective-action Output Field Descriptions**

| Label                                         | Description                                                                                                                                   |
|-----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| Secondary Action                              | The status of the target of the secondary action, if configured, based on extended checks.                                                    |
| Downloaded Action                             | The action downloaded by the CPM to the IOM.                                                                                                  |
| Stickiness Timer                              | The status of the stickiness timer, if any.                                                                                                   |
| Effective Action based on application context | Section of the output providing the effective action, in the context of services, that a packet matching the criteria will be subject to.     |
| Service Id                                    | The service ID on which the filter policy ID is applied. The output also provides a list of service points where the filter has been applied. |
| Type                                          | The service type in which the service has been applied.                                                                                       |
| Ingress/Egress                                | The direction in which the service has been applied.                                                                                          |
| Effective Action                              | Indicates the effective action the packet will be subject to.                                                                                 |
| Extended Action                               | Indicates the configured extended action, if any.                                                                                             |

## log

**Syntax** **log** *log-id* [**match** *string*]  
**log** [**bindings**]

**Context** show>filter

**Description** This command shows the contents of a memory-based or a file-based filter log.

If the optional keyword **match** and *string* parameter are given, the command displays the given filter log from the first occurrence of the given string.

**Parameters** *log-id* — the filter log ID destination expressed as a decimal integer

**Values** 101 to 199

**match** *string* — specifies to start displaying the filter log entries from the first occurrence of *string*

**bindings** — displays the number of filter logs currently instantiated

**Output** The following output is an example of filter log entry information, and [Table 64](#) describes the fields. If log summary is active, the filter log mini-tables contain the information described in [Table 65](#).

### Sample Output

```

2007/04/13 16:23:09 Filter: 100:100 Desc: Entry-100
Interface: to-ser1 Action: Forward
Src MAC: 04-5b-01-01-00-02 Dst MAC: 04-5d-01-01-00-02 EtherType: 0800
Src IP: 10.10.0.1:646 Dst IP: 10.10.0.4:49509 Flags: TOS: c0
Protocol: TCP Flags: ACK

```

```

2007/04/13 16:23:10 Filter: 100:100 Desc: Entry-100
Interface: to-ser1 Action: Forward
Src MAC: 04-5b-01-01-00-02 Dst MAC: 04-5d-01-01-00-02 EtherType: 0800
Src IP: 10.10.0.1:646 Dst IP: 10.10.0.3:646 Flags: TOS: c0
Protocol: UDP

```

```

2007/04/13 16:23:12 Filter: 100:100 Desc: Entry-100
Interface: to-ser1 Action: Forward
Src MAC: 04-5b-01-01-00-02 Dst MAC: 01-00-5e-00-00-05 EtherType: 0800
Src IP: 10.10.13.1 Dst IP: 224.0.0.5 Flags: TOS: c0
Protocol: 89
Hex: 02 01 00 30 0a 0a 00 01 00 00 00 00 00 ba 90 00 00
 00 00 00 00 00 00 00 00 ff ff ff 00 00 03 02 01

```

```
A:ALA-A>config# show filter log bindings
```

```
=====
Filter Log Bindings
=====
```

```

Total Log Instances (Allowed) : 2046
Total Log Instances (In Use) : 0
Total Log Bindings : 0

```

```

Type FilterId EntryId Log Instantiated

```

```
No Instances found
```

```
=====
A:ALA-A>config#
```

A summary log will be printed only in case TotCnt is different from 0. Only the address types with at least 1 entry in the minitable will be printed.

```
A:ALA-A>config# show filter log 190
```

```
=====
Summary Log[190] Crit1: SrcAddr TotCnt: 723 ArpCnt: 83
```

```

Mac 8 06-06-06-06-06-06
Mac 8 06-06-06-06-06-05
Mac 8 06-06-06-06-06-04
Mac 8 06-06-06-06-06-03
Mac 8 06-06-06-06-06-02
Ip 16 6.6.6.1
Ip 16 6.6.6.2
Ip 16 6.6.6.3
Ip 16 6.6.6.4
Ip 8 6.6.6.5
Ipv6 8 3FE:1616:1616:1616:1616:1616::
Ipv6 8 3FE:1616:1616:1616:1616:1616:FFFF:FFFF
Ipv6 8 3FE:1616:1616:1616:1616:1616:FFFF:FFFE
Ipv6 8 3FE:1616:1616:1616:1616:1616:FFFF:FFFD
Ipv6 8 3FE:1616:1616:1616:1616:1616:FFFF:FFFC

```

```
=====
A:ALA-A
```



**Log Message Formatting** — Each filter log entry contains the following information in case summary log feature is not active (as appropriate).

**Table 64 Filter Log Output Fields**

| Label                          | Description                                                                                                                                                                                              |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>yyyy/mm/dd<br/>hh:mm:ss</i> | The date and timestamp for the log filter entry where <i>yyyy</i> is the year, <i>mm</i> is the month, <i>dd</i> is the day, <i>hh</i> is the hour, <i>mm</i> is the minute and <i>ss</i> is the second. |
| Filter                         | The filter ID and the entry ID which generated the filter log entry in the form <i>Filter_ID:Entry_ID</i> .                                                                                              |
| Desc                           | The description of the filter entry ID which generated the filter log entry.                                                                                                                             |
| Interface                      | The IP interface on which the filter ID and entry ID was associated which generated the filter log entry.                                                                                                |
| Action                         | The action of the filter entry on the logged packet.                                                                                                                                                     |
| Src MAC                        | The source MAC address of the logged packet.                                                                                                                                                             |
| Dst MAC                        | The destination MAC of the logged packet.                                                                                                                                                                |
| EtherType                      | The Ethernet type of the logged Ethernet type II packet.                                                                                                                                                 |
| Src IP                         | The source IP address of the logged packet. The source port will be displayed after the IP address as appropriate separated with a colon.                                                                |
| Dst IP                         | The destination IP address of the logged packet. The source port will be displayed after the IP address as appropriate separated with a colon.                                                           |
| Flags<br>(IP flags)            | M — The more fragments IP flag is set in the logged packet.<br>DF — The do not fragment IP flag is set in the logged packet.                                                                             |
| TOS                            | The TOS byte value in the logged packet.                                                                                                                                                                 |
| Protocol                       | The IP protocol of the logged packet (TCP, UDP, ICMP or a protocol number in hex).                                                                                                                       |
| Flags<br>(TCP flags)           | URG — Urgent bit set.<br>ACK — Acknowledgment bit set.<br>RST — Reset bit set.<br>SYN — Synchronize bit set.<br>FIN — Finish bit set.                                                                    |

**Table 64** Filter Log Output Fields (Continued)

| Label                         | Description                                                                                                                                                                                                                                                                     |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| HEX                           | If an IP protocol does not have a supported decode, the first 32 bytes following the IP header are printed in a hex dump.<br>Log entries for non-IP packets include the Ethernet frame information and a hex dump of the first 40 bytes of the frame after the Ethernet header. |
| Total Log Instances (Allowed) | Specifies the maximum allowed instances of filter logs allowed on the system.                                                                                                                                                                                                   |
| Total Log Instances (In Use)  | Specifies the instances of filter logs presently existing on the system.                                                                                                                                                                                                        |
| Total Log Bindings            | Specifies the count of the filter log bindings presently existing on the system.                                                                                                                                                                                                |
| Type                          | The type of service of the service ID.                                                                                                                                                                                                                                          |
| Filter ID                     | Uniquely identifies an IP filter as configured on the system.                                                                                                                                                                                                                   |
| Entry ID                      | The identifier which uniquely identifies an entry in a filter table.                                                                                                                                                                                                            |
| Log                           | Specifies an entry in the filter log table.                                                                                                                                                                                                                                     |
| Instantiated                  | Specifies if the filter log for this filter entry has or has not been instantiated.                                                                                                                                                                                             |

If the packet being logged does not have a source or destination MAC address (that is, POS) then the MAC information output line is omitted from the log entry.

If log summary is active, the filter log mini-tables contain the information described in [Table 65](#).

**Table 65** Filter Log Summary Mini-Table Fields

| Label                    | Description                                                              |
|--------------------------|--------------------------------------------------------------------------|
| <i>Summary Log LogID</i> | Displays the log ID.                                                     |
| Crit1                    | Summary criterion that is used as index into the mini-tables of the log. |
| TotCnt                   | The total count of logs.                                                 |
| ArpCnt                   | Displays the total number of ARP messages logged for this log ID.        |
| Src...<br>Dst...         | The address type indication of the key in the mini-table.                |

**Table 65 Filter Log Summary Mini-Table Fields (Continued)**

| Label   | Description                                                                  |
|---------|------------------------------------------------------------------------------|
| count   | The number of messages logged with the specified source/destination address. |
| address | The address for which count messages were received.                          |

## mac

**Syntax** **mac** *mac-filter-id*  
**mac** *mac-filter-id* **associations**  
**mac** *mac-filter-id* [**type** *entry-type*] **counters** [**detail**]  
**mac** [*mac-filter-id*] **embedded** [**inactive**]  
**mac** *mac-filter-id* **entry** *entry-id* [**counters**] [**detail**]  
**mac** [**filter-type** *filter-type*]  
**mac** *mac-filter-id* **type** *entry-type*  
**mac** *mac-filter-id* [**entry** *entry-id*] **effective-action** [**router** | **service** {*service-id* | *service-name*}] [**ingress** | **egress**]

**Context** show>filter

**Description** This command displays MAC filter information.

When **effective-action** is specified, this command displays what effectively happens to a packet that matches the criteria associated with the MAC filter policy.

**Parameters** *mac-filter-id* — Displays detailed information for the specified filter ID and its filter entries.

**Values** 1 to 65535

**associations** — Appends information as to where the filter policy ID is applied to the detailed filter policy ID output.

**counters** — Displays counter information for the specified filter ID.

**entry** *entry-id* — Displays information on the specified filter entry ID for the specified filter ID only.

**Values** 1 to 65535

**type** *entry-type* — Specifies the type of filter entries as “fixed” or “embedded”.

**filter-type** *filter-type* — Specifies the type of filter entries as “config” or “vsd”.

**effective-action** — Displays the action that the system will effectively apply to the packet.

**router** — Filters the output and only displays the information for that specific service (“Base” instance).

**service service-id** — Filters the output and only displays the information for the specified service. The specified value must correspond to an existing service in which the filter has been applied.

**service service-name** — Filters the output and only displays the information for the specified service. The specified value must correspond to an existing service in which the filter has been applied.

**ingress** — Filters the output and only displays the information for filter policies applied on ingress.

**egress** — Filters the output and only displays the information for filter policies applied on egress.

**Output**     **No Parameters Specified** — A brief listing of MAC filters is produced when no parameters are specified; [Table 66](#) describes the output fields.

### Sample Output

```
*A:Dut-C# show filter mac
=====
Configured Mac Filters Total: 1
=====
Filter-Id Scope Applied Description Type

10 Template No normal
=====
Num Mac filters: 1
=====
*A:Dut-C#
```

**Table 66**     **Filter MAC Output Fields**

| Label       | Description                                          |
|-------------|------------------------------------------------------|
| Filter Id   | The MAC filter ID                                    |
| Scope       | Template<br>The filter policy is of type Template.   |
|             | Exclusive<br>The filter policy is of type Exclusive. |
| Applied     | No<br>The filter policy ID has not been applied.     |
|             | Yes<br>The filter policy ID is applied.              |
| Description | The MAC filter policy description.                   |

**Filter ID Specified** — The following output is an example of MAC filter information when the filter ID is specified, and [Table 67](#) describes the fields. Detailed filter information for the filter ID and its entries is produced when the filter ID is specified.

### Sample Output

```
=====
Mac Filter : 200
=====
Filter Id : 200 Applied : No
Scope : Exclusive D. Action : Drop
Description : Forward SERVER sourced packets

Filter Match Criteria : Mac

Entry : 200 FrameType : 802.2SNAP
Description : Not Available
Src Mac : 00:00:5a:00:00:00 ff:ff:ff:00:00:00
Dest Mac : 00:00:00:00:00:00 00:00:00:00:00:00
Dot1p : Undefined Ethertype : 802.2SNAP
DSAP : Undefined SSAP : Undefined
Snap-pid : Undefined ESnap-oui-zero : Undefined
Match action : Forward
Ing. Matches : 0 Egr. Matches : 0
Entry : 300 (Inactive) FrameType : Ethernet
Description : Not Available
Src Mac : 00:00:00:00:00:00 00:00:00:00:00:00
Dest Mac : 00:00:00:00:00:00 00:00:00:00:00:00
Dot1p : Undefined Ethertype : Ethernet
DSAP : Undefined SSAP : Undefined
Snap-pid : Undefined ESnap-oui-zero : Undefined
Match action : Default
Ing. Matches : 0 Egr. Matches : 0
=====
```

**Table 67** Filter MAC with Filter-ID Specified Output Fields

| Label                   | Description                                          |
|-------------------------|------------------------------------------------------|
| MAC Filter<br>Filter Id | The MAC filter policy ID.                            |
| Scope                   | Template<br>The filter policy is of type Template.   |
|                         | Exclusive<br>The filter policy is of type Exclusive. |
| Description             | The MAC filter policy description.                   |

**Table 67 Filter MAC with Filter-ID Specified Output Fields (Continued)**

| Label                 | Description                                                                                                                                          |
|-----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| Applied               | No<br>The filter policy ID has not been applied.                                                                                                     |
|                       | Yes<br>The filter policy ID is applied.                                                                                                              |
| Def. Action           | Forward<br>The default action for the filter ID for packets that do not match the filter entries is to forward.                                      |
|                       | Drop<br>The default action for the filter ID for packets that do not match the filter entries is to drop.                                            |
| Filter Match Criteria | MAC<br>Indicates the filter is an MAC filter policy.                                                                                                 |
| Entry                 | The filter entry ID. If the filter entry ID indicates the entry is (Inactive), then the filter entry is incomplete as no action has been specified.  |
| Description           | The filter entry description.                                                                                                                        |
| FrameType             | Ethernet<br>The entry ID match frame type is Ethernet IEEE 802.3.                                                                                    |
|                       | Ethernet II<br>The entry ID match frame type is Ethernet Type II.                                                                                    |
| Src MAC               | The source MAC address and mask match criterion. When both the MAC address and mask are all zeros, no criterion specified for the filter entry.      |
| Dest MAC              | The destination MAC address and mask match criterion. When both the MAC address and mask are all zeros, no criterion specified for the filter entry. |
| Dot1p                 | The IEEE 802.1p value for the match criteria. Undefined indicates no value is specified.                                                             |
| Ethertype             | The Ethertype value match criterion.                                                                                                                 |
| DSAP                  | The DSAP value match criterion.<br>Undefined indicates no value specified.                                                                           |
| SSAP                  | SSAP value match criterion. Undefined indicates no value specified.                                                                                  |
| Snap-pid              | The Ethernet SNAP PID value match criterion. Undefined indicates no value specified.                                                                 |

**Table 67 Filter MAC with Filter-ID Specified Output Fields (Continued)**

| Label          | Description                                                                                                                                                                                              |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Esnap-oui-zero | Non-Zero<br>Filter entry matches a non-zero value for the Ethernet SNAP OUI.                                                                                                                             |
|                | Zero<br>Filter entry matches a zero value for the Ethernet SNAP OUI.                                                                                                                                     |
|                | Undefined<br>No Ethernet SNAP OUI value specified.                                                                                                                                                       |
| Match action   | Default<br>The filter does not have an explicit forward or drop match action specified. If the filter entry ID indicates the entry is inactive, the filter entry is incomplete, no action was specified. |
|                | Drop<br>Packets matching the filter entry criteria will be dropped.                                                                                                                                      |
|                | Forward<br>Packets matching the filter entry criteria is forwarded.                                                                                                                                      |
| Ing. Matches   | The number of ingress filter matches/hits for the filter entry.                                                                                                                                          |
| Egr. Matches   | The number of egress filter matches/hits for the filter entry.                                                                                                                                           |

**Filter Associations** — The associations for a filter ID will be displayed if the **associations** keyword is specified. The association information is appended to the filter information.

The following output is an example of MAC filter information when the associations keyword is specified, and [Table 68](#) describes the fields.

### Sample Output

```
A:ALA-49# show filter mac 3 associations
=====
Mac Filter
=====
Filter ID : 3 Applied : Yes
Scope : Template Def. Action : Drop
Entries : 1

Filter Association : Mac

Service Id : 1001 Type : VPLS
- SAP 1/1/1:1001 (Egress)
=====
```

**Table 68** Filter MAC Associations Output Fields

| Label              | Description                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| Filter Association | Mac<br>The filter associations displayed are for a MAC filter policy ID.                                                                      |
| Service Id         | The service ID on which the filter policy ID is applied. The output also provides a list of service points where the filter has been applied. |
| SAP                | The Service Access Point or spoke/mesh SDP on which the filter policy ID is applied.                                                          |
| Type               | The type of service of the Service ID.                                                                                                        |
| (Ingress)          | The filter policy ID is applied as an ingress filter policy on the interface.                                                                 |
| (Egress)           | The filter policy ID is applied as an egress filter policy on the interface.                                                                  |

**Filter Entry Counters Output** — When the **counters** keyword is specified, the filter entry output displays the filter matches/hit information. The following table describes the command output for the command.

```
A:ALA-49# show filter mac 8 counters
=====
Mac Filter
=====
Filter Id : 8 Applied : Yes
Scope : Template Def. Action : Forward
Entries : 2
Description : Description for Mac Filter Policy id # 8

Filter Match Criteria : Mac

Entry : 8 FrameType : Ethernet
Ing. Matches: 80 pkts (5440 bytes)
Egr. Matches: 62 pkts (3968 bytes)

Entry : 10 FrameType : Ethernet
Ing. Matches: 80 pkts (5440 bytes)
Egr. Matches: 80 pkts (5120 bytes)
=====
```

**Table 69** Filter MAC Counters Output Field Descriptions

| Label                   | Description               |
|-------------------------|---------------------------|
| Mac Filter<br>Filter Id | The MAC filter policy ID. |



**Table 69 Filter MAC Counters Output Field Descriptions (Continued)**

| Label                 | Description                                                                                                                                         |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| Scope                 | Template<br>The filter policy is of type Template.                                                                                                  |
|                       | Exclusive<br>The filter policy is of type Exclusive.                                                                                                |
| Description           | The MAC filter policy description.                                                                                                                  |
| Applied               | No<br>The filter policy ID has not been applied.                                                                                                    |
|                       | Yes<br>The filter policy ID is applied.                                                                                                             |
| Def. Action           | Forward<br>The default action for the filter ID for packets that do not match the filter entries is to forward.                                     |
|                       | Drop<br>The default action for the filter ID for packets that do not match the filter entries is to drop.                                           |
| Filter Match Criteria | Mac<br>Indicates the filter is an MAC filter policy.                                                                                                |
| Entry                 | The filter entry ID. If the filter entry ID indicates the entry is (Inactive), then the filter entry is incomplete as no action has been specified. |
| FrameType             | Ethernet<br>The entry ID match frame type is Ethernet IEEE 802.3.                                                                                   |
|                       | 802.2LLC<br>The entry ID match frame type is Ethernet IEEE 802.2 LLC.                                                                               |
|                       | 802.2SNAP<br>The entry ID match frame type is Ethernet IEEE 802.2 SNAP.                                                                             |
|                       | Ethernet II<br>The entry ID match frame type is Ethernet Type II.                                                                                   |
| Ing. Matches          | The number of ingress filter matches/hits for the filter entry.                                                                                     |
| Egr. Matches          | The number of egress filter matches/hits for the filter entry.                                                                                      |

**Show Filter MAC Output (with effective-action specified)** — The following is a sample output of MAC filter information when the **effective-action** keyword is specified. [Table 70](#) describes the command output fields.

If the main action (either primary or secondary) cannot be performed, a reason will be given. This will be displayed on the same line as the Effective Action. The reason codes as currently defined are:

- entry-default
- filter-default-action
- pbr-down-action-override
- action not supported on egress
- target does not exist

### Sample Output

```
show filter mac 1 effective-action
=====
Mac Filter
=====
Filter Id : 1 Applied : Yes
Scope : Template Def. Action : Drop
Entries : 1 Type : normal
Description : (Not Specified)

Entry : 2

Stickiness : No
PBR Dwn Act Override: None
PBR Down Action : Drop (entry-default)

Configuration
Primary Action : Forward (SAP)
 Next Hop : 1/1/3
 Service Id : Not configured
Secondary Action : None

Status
Target status based on extended checks
 Primary Action : Target does not exist
 Secondary Action : None
Downloaded Action : Forward
Stickiness Timer : Not Running

Effective Action based on application context
Service Id : 10 Type : VPLS
Ingress
 Effective Action: Forward (target does not exist)
=====
```

**Table 70 Show Filter MAC effective-action Output Field Descriptions**

| Label                | Description                                                                                                                                         |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------|
| Filter Id            | The MAC filter policy ID.                                                                                                                           |
| Applied              | No<br>The filter policy ID has not been applied.                                                                                                    |
|                      | Yes<br>The filter policy ID is applied.                                                                                                             |
| Scope                | Template<br>The filter policy is of type Template.                                                                                                  |
|                      | Exclusive<br>The filter policy is of type Exclusive.                                                                                                |
| Def. Action          | Forward<br>The default action for the filter ID for packets that do not match the filter entries is to forward.                                     |
|                      | Drop<br>The default action for the filter ID for packets that do not match the filter entries is to drop.                                           |
| Entries              | The number of entries configured in this filter ID.                                                                                                 |
| Type                 | The type of entries configured in this filter.                                                                                                      |
| Description          | The MAC filter policy description.                                                                                                                  |
| Entry                | The filter entry ID. If the filter entry ID indicates the entry is (Inactive), then the filter entry is incomplete as no action has been specified. |
| Stickiness           | No<br>Stickiness is not configured.                                                                                                                 |
|                      | Yes<br>Stickiness is configured.                                                                                                                    |
| PBR Dwn Act Override | Indicates whether or not the action to take when the PBR target is down has been overridden.                                                        |
| PBR Down Action      | The action to take in case the target is down.                                                                                                      |
| Configuration        | Section of the output providing information on the configured parameters.                                                                           |
| Primary Action       | The configured action, if any. Indented sub-labels in the show output provide configured parameters for this action.                                |

**Table 70** Show Filter MAC effective-action Output Field Descriptions

| Label                                         | Description                                                                                                                                   |
|-----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| Secondary Action                              | The configured secondary action, if any. Indented sub-labels in the show output provide configured parameters for this action.                |
| Status/Target status based on extended checks | Section of the output providing information on the operational status of certain parameters.                                                  |
| Primary Action                                | The status of the target of the primary action, if configured, based on extended checks.                                                      |
| Secondary Action                              | The status of the target of the secondary action, if configured, based on extended checks.                                                    |
| Downloaded Action                             | The action downloaded by the CPM to the IOM.                                                                                                  |
| Stickiness Timer                              | The status of the stickiness timer, if any.                                                                                                   |
| Effective Action based on application context | Section of the output providing the effective action, in the context of services, that a packet matching the criteria will be subject to.     |
| Service Id                                    | The service ID on which the filter policy ID is applied. The output also provides a list of service points where the filter has been applied. |
| Type                                          | The service type in which the service has been applied.                                                                                       |
| Ingress/Egress                                | The direction in which the service has been applied.                                                                                          |
| Effective Action                              | Indicates the effective action the packet will be subject to.                                                                                 |

## redirect-policy

**Syntax** **redirect-policy** [*redirect-policy-name* {**dest** *ip-address*} [**associations**]]

**Context** show>filter

**Description** This command shows redirect filter information.

**Parameters** *redirect-policy-name* — displays information for the specified redirect policy  
**dest** *ip-address* — directs the router to use a specified IP address for communication  
**associations** — appends association information

**Output** **Redirect Policy Output** — The following output is an example of redirect policy information, and [Table 71](#) describes the fields.

### Sample Output

```
A:ALA-A>config>filter# show filter redirect-policy
=====
Redirect Policies
=====
Redirect Policy Applied Description

wccp Yes
redirect1 Yes New redirect info
redirect2 Yes Test test test test
=====
ALA-A>config>filter#

ALA-A>config>filter# show filter redirect-policy redirect1
=====
Redirect Policy
=====
Redirect Policy: redirect1 Applied : Yes
Description : New redirect info
Active Dest : 10.10.10.104

Destination : 10.10.10.104

Description : SNMP_to_104
Admin Priority : 105 Oper Priority: 105
Admin State : Up Oper State : Up

SNMP Test : SNMP-1
Interval : 30 Timeout : 1
Drop Count : 30
Hold Down : 120 Hold Remain : 0
Last Action at : None Taken

Destination : 10.10.10.105

Description : another test
Admin Priority : 95 Oper Priority: 105
Admin State : Up Oper State : Down

Ping Test
Interval : 1 Timeout : 30
Drop Count : 5
Hold Down : 0 Hold Remain : 0
Last Action at : 03/19/2007 00:46:55 Action Taken : Disable

Destination : 10.10.10.106

Description : (Not Specified)
Admin Priority : 90 Oper Priority: 90
Admin State : Up Oper State : Down

URL Test : URL_to_Proxy
Interval : 10 Timeout : 10
Drop Count : 3
Hold Down : 0 Hold Remain : 0
Last Action at : 03/19/2007 05:04:15 Action Taken : Disable
Priority Change : 0 Return Code : 0
=====
A:ALA-A>config>filter#
```

```

A:ALA-A>show filter redirect-policy redirect1 dest 10.10.10.106
=====
Redirect Policy
=====
Redirect Policy: redirect1 Applied : Yes
Description : New redirect info
Active Dest : 10.10.10.104

Destination : 10.10.10.106

Description : (Not Specified)
Admin Priority : 90 Oper Priority: 90
Admin State : Up Oper State : Down

URL Test : URL_to_Proxy
Interval : 10 Timeout : 10
Drop Count : 3
Hold Down : 0 Hold Remain : 0
Last Action at : 03/19/2007 05:04:15 Action Taken : Disable
Priority Change: 0 Return Code : 0
=====
ALA-A#

```

**Table 71** Filter Redirect-Policy Output Fields

| Label              | Description                                                                                                                                                               |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Redirect Policy    | Specifies a specific redirect policy.                                                                                                                                     |
| Applied            | Specifies whether the redirect policy is applied to a filter policy entry.                                                                                                |
| Description        | Displays the user-provided description for this redirect policy.                                                                                                          |
| Active Destination | ip address<br>Specifies the IP address of the active destination.                                                                                                         |
|                    | none<br>Indicates that there is currently no active destination.                                                                                                          |
| Destination        | Specifies the destination IP address.                                                                                                                                     |
| Oper Priority      | Specifies the operational value of the priority for this destination. The highest operational priority across multiple destinations is used as the preferred destination. |
| Admin Priority     | Specifies the configured base priority for the destination.                                                                                                               |
| Admin State        | Specifies the configured state of the destination.                                                                                                                        |
|                    | Out of Service<br>Tests for this destination will not be conducted.                                                                                                       |

**Table 71 Filter Redirect-Policy Output Fields (Continued)**

| Label          | Description                                                                                                                                                                                      |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Oper State     | Specifies the operational state of the destination.                                                                                                                                              |
| Ping Test      | Specifies the name of the ping test.                                                                                                                                                             |
| Timeout        | Specifies the amount of time in seconds that is allowed for receiving a response from the far-end host. If a reply is not received within this time the far-end host is considered unresponsive. |
| Interval       | Specifies the amount of time in seconds between consecutive requests sent to the far end host.                                                                                                   |
| Drop Count     | Specifies the number of consecutive requests that must fail for the destination to declared unreachable.                                                                                         |
| Hold Down      | Specifies the amount of time in seconds that the system should be held down if any of the test has marked it unreachable.                                                                        |
| Hold Remain    | Specifies the amount of time in seconds that the system will remain in a hold down state before being used again.                                                                                |
| Last Action at | Displays a time stamp of when this test received a response for a probe that was sent out.                                                                                                       |
| SNMP Test      | Specifies the name of the SNMP test.                                                                                                                                                             |
| URL Test       | Specifies the name of the URL test.                                                                                                                                                              |

## system-filter

**Syntax** **system-filter [chained-to]**

**Context** show>filter

**Description** This command shows system filter information.

**Parameters** **chained-to** — this option displays filters that chain to a given system filter

**Output** **No Parameters Specified** — When no parameters are specified, the output is grouped for IPv4 and IPv6, and displays information about the active system filter and all filters with scope **system**.

The following output is an example of system filter information when no parameters are specified.

### Sample Output

```
*A:Dut-C>show>filter# system-filter
```

```
=====
IP system filters
=====
Filter-Id Active

100 Yes
65535 No

No. of IP system filters (total / active): 2 / 1
=====

=====
IPv6 system filters
=====
Filter-Id Active

No Matching Entries

No. of IPv6 system filters (total / active): 0 / 0
=====
```

**With chained-to Option Specified** — The following output is an example of system filter information when the **chained-to** option is specified, .

```
*A:Dut-C>show>filter# system-filter chained-to

=====
IP filters that chain to the active IP system filter
=====
3 4 5 6
5:23 6:24

No. of IP filters that chain to the active IP system filter: 6
=====

=====
IPv6 filters that chain to the active IPv6 system filter
=====
No Matching Entries

No. of IPv6 filters that chain to the active IPv6 system filter: 0
=====
```

match-list

|             |                                                                                                                   |
|-------------|-------------------------------------------------------------------------------------------------------------------|
| Syntax      | match-list                                                                                                        |
| Context     | show>filter                                                                                                       |
| Description | This command enables the context to display information for match lists used in filter policies (IOM/FP and CPM). |



## ip-prefix-list

|                    |                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ip-prefix-list</b> [ <i>prefix-list-name</i> ]<br><b>ip-prefix-list</b> <i>prefix-list-name</i> <b>references</b>                                                       |
| <b>Context</b>     | show>filter>match-list                                                                                                                                                     |
| <b>Description</b> | This command displays IPv4 prefixes information for match criteria in IPv4 ACL and CPM filter policies.                                                                    |
| <b>Parameters</b>  | <i>prefix-list-name</i> — a string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes. |

## ipv6-prefix-list

|                    |                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ipv6-prefix-list</b> [ <i>prefix-list-name</i> ]<br><b>ipv6-prefix-list</b> <i>prefix-list-name</i> <b>references</b>                                                   |
| <b>Context</b>     | show>filter>match-list                                                                                                                                                     |
| <b>Description</b> | This command displays IPv6 prefixes information for match criteria in IPv6 ACL and CPM filter policies.                                                                    |
| <b>Parameters</b>  | <i>prefix-list-name</i> — a string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes. |

## port-list

|                    |                                                                                                                                                                          |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>port-list</b> [ <i>port-list-name</i> ]<br><b>port-list</b> <i>port-list-name</i> <b>references</b>                                                                   |
| <b>Context</b>     | show>filter>match-list                                                                                                                                                   |
| <b>Description</b> | This command displays TCP/UDP/SCTP port values or ranges for match criteria in IPv4 and IPv6 ACL and CPM filter policies.                                                |
| <b>Parameters</b>  | <i>port-list-name</i> — a string of up to 32 characters of printable ASCII characters. If special characters are used, the string must be enclosed within double quotes. |

### 4.5.2.2 Clear Commands

## ip

|               |                                                                                                |
|---------------|------------------------------------------------------------------------------------------------|
| <b>Syntax</b> | <b>ip</b> <i>filter-id</i> [ <b>entry</b> <i>entry-id</i> ] [ <b>ingress</b>   <b>egress</b> ] |
|---------------|------------------------------------------------------------------------------------------------|

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | clear>filter                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | <p>Clears the counters associated with the entries of the specified IPv4 filter policy.</p> <p>By default, the counters associated with each entry of the specified filter policy are all cleared. The scope of which counters are cleared can be narrowed using the command line parameters.</p>                                                                                                                                                                                                                                     |
| <b>Default</b>     | Clears all counters associated with each entry of the specified IPv4 filter policy.                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b>  | <p><i>filter-id</i> — the IPv4 filter policy ID for which to clear the entry counters. Values can be expressed in different formats. The following shows decimal integer format.</p> <p><b>Values</b> 1 to 65535</p> <p><b>entry</b> <i>entry-id</i> — specifies that only the counters associated with the specified filter policy entry will be cleared</p> <p><b>Values</b> 1 to 65535</p> <p><b>ingress</b> — specifies to only clear the ingress counters</p> <p><b>egress</b> — specifies to only clear the egress counters</p> |

## ipv6

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ipv6</b> <i>ipv6-filter-id</i> [ <b>entry</b> <i>entry-id</i> ] [ <b>ingress</b>   <b>egress</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | clear>filter                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b> | <p>Clears the counters associated with the entries of the specified IPv6 filter policy.</p> <p>By default, the counters associated with each entry of the specified filter policy are all cleared. The scope of which counters are cleared can be narrowed using the command line parameters.</p>                                                                                                                                                                                                                                               |
| <b>Default</b>     | Clears all counters associated with each entry of the specified IPv6 filter policy.                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Parameters</b>  | <p><i>ipv6-filter-id</i> — the IPv6 filter policy ID for which to clear the entry counters. Values can be expressed in different formats. The following only shows decimal integer format.</p> <p><b>Values</b> 1 to 65535</p> <p><b>entry</b> <i>entry-id</i> — specifies that only the counters associated with the specified filter policy entry will be cleared</p> <p><b>Values</b> 1 to 65535</p> <p><b>ingress</b> — specifies to only clear the ingress counters</p> <p><b>egress</b> — specifies to only clear the egress counters</p> |

## log

|                    |                                                                                                                           |
|--------------------|---------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>log</b> <i>log-id</i>                                                                                                  |
| <b>Context</b>     | clear>filter                                                                                                              |
| <b>Description</b> | Clears the contents of a memory or file based filter log.<br><br>This command has no effect on a syslog based filter log. |
| <b>Parameters</b>  | <i>log-id</i> — the filter log ID expressed as a decimal integer                                                          |
| <b>Values</b>      | 101 to 199                                                                                                                |

## mac

|                    |                                                                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>mac</b> <i>mac-filter-id</i> [ <b>entry</b> <i>entry-id</i> ] [ <b>ingress</b>   <b>egress</b> ]                                                                                                                                                                                       |
| <b>Context</b>     | clear>filter                                                                                                                                                                                                                                                                              |
| <b>Description</b> | Clears the counters associated with the entries of the specified MAC filter policy.<br><br>By default, the counters associated with each entry of the specified filter policy are all cleared. The scope of which counters are cleared can be narrowed using the command line parameters. |
| <b>Default</b>     | Clears all counters associated with each entry of the specified MAC filter policy.                                                                                                                                                                                                        |
| <b>Parameters</b>  | <i>mac-filter-id</i> — the MAC filter policy ID for which to clear the entry counters. Values can either be expressed as a decimal integer or as an ASCII string of up to 64 characters. The following values only shows decimal integer.                                                 |
| <b>Values</b>      | 1 to 65535                                                                                                                                                                                                                                                                                |
|                    | <b>entry</b> <i>entry-id</i> — specifies that only the counters associated with the specified filter policy entry will be cleared. The values are expressed as a decimal integer.                                                                                                         |
| <b>Values</b>      | 1 to 65535                                                                                                                                                                                                                                                                                |
|                    | <b>ingress</b> — specifies to only clear the ingress counters                                                                                                                                                                                                                             |
|                    | <b>egress</b> — specifies to only clear the egress counters                                                                                                                                                                                                                               |

### 4.5.2.3 Monitor Commands

The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

---

## ip

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ip</b> <i>ip-filter-id</i> <b>entry</b> <i>entry-id</i> [ <b>interval</b> <i>seconds</i> ] [ <b>repeat</b> <i>repeat</i> ] [ <b>absolute</b>   <b>rate</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Context</b>     | monitor>filter                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | This command monitors the counters associated with the specified entry of the specified IP filter policy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Parameters</b>  | <p><i>filter-id</i> — the IPv4 filter policy ID. Values can be expressed in different formats. The following only shows decimal integer format values.</p> <p><b>Values</b> 1 to 65535</p> <p><b>entry</b> <i>entry-id</i> — specifies the filter policy entry to monitor, as a decimal integer</p> <p><b>Values</b> 1 to 65535</p> <p><b>interval</b> <i>seconds</i> — configures the interval for each display in seconds</p> <p><b>Default</b> 10 seconds</p> <p><b>Values</b> 3 to 60</p> <p><b>repeat</b> <i>repeat</i> — configures how many times the command is repeated</p> <p><b>Default</b> 10</p> <p><b>Values</b> 1 to 999</p> <p><b>absolute</b> — when the <b>absolute</b> keyword is specified, the raw statistics are displayed, without processing. No calculations are performed on the delta or rate statistics.</p> <p><b>rate</b> — when the <b>rate</b> keyword is specified, the rate-per-second for each statistic is displayed instead of the delta</p> <p><b>Default</b> <b>absolute</b></p> |

## ipv6

|                    |                                                                                                                                                                                                                                                                                                                                          |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ipv6</b> <i>ipv6-filter-id</i> <b>entry</b> <i>entry-id</i> [ <b>interval</b> <i>seconds</i> ] [ <b>repeat</b> <i>repeat</i> ] [ <b>absolute</b>   <b>rate</b> ]                                                                                                                                                                      |
| <b>Context</b>     | monitor>filter                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | This command monitors the counters associated with the IPv6 filter policy.                                                                                                                                                                                                                                                               |
| <b>Parameters</b>  | <p><i>ipv6-filter-id</i> — the IPv6 filter policy ID. Values can be expressed in different formats. The following only shows decimal integer format values.</p> <p><b>Values</b> 1 to 65535</p> <p><b>entry</b> <i>entry-id</i> — specifies the filter policy entry to monitor, as a decimal integer</p> <p><b>Values</b> 1 to 65535</p> |

**interval** *seconds* — configures the interval for each display in seconds

**Default** 10 seconds

**Values** 3 to 60

**repeat** *repeat* — configures how many times the command is repeated

**Default** 10

**Values** 1 to 999

**absolute** — when the **absolute** keyword is specified, the raw statistics are displayed, without processing. No calculations are performed on the delta or rate statistics.

**rate** — when the **rate** keyword is specified, the rate-per-second for each statistic is displayed instead of the delta

**Default** **absolute**

## mac

**Syntax** **mac** *mac-filter-id* **entry** *entry-id* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**]

**Context** monitor>filter

**Description** This command monitors the counters associated with the specified entry of the specified MAC filter policy.

**Parameters** *mac-filter-id* — the MAC filter policy ID. Values can be expressed in different formats. The following only shows decimal integer format values.

**Values** 1 to 65535

**entry** *entry-id* — specifies the filter policy entry to monitor, as a decimal integer

**Values** 1 to 65535

**interval** *seconds* — configures the interval for each display in seconds

**Default** 10 seconds

**Values** 3 to 60

**repeat** *repeat* — configures how many times the command is repeated

**Default** 10

**Values** 1 to 999

**absolute** — when the **absolute** keyword is specified, the raw statistics are displayed, without processing. No calculations are performed on the delta or rate statistics.

**rate** — when the **rate** keyword is specified, the rate-per-second for each statistic is displayed instead of the delta

**Default** **absolute**

### 4.5.2.4 Debug Commands

The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

#### cpm

|                    |                                                                                                                                                                                     |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>cpm</b>                                                                                                                                                                          |
| <b>Context</b>     | tools>dump>filter>resources                                                                                                                                                         |
| <b>Description</b> | This command displays information about filter resource utilization on the CPM, consumption by filter-using services such as OpenFlow, and the filters that use the most resources. |
| <b>Output</b>      | The following output is an example of filter resource utilization information.                                                                                                      |

#### Sample Output

```
*A:Dut-C>tools>dump>filter>resources># cpm

=====
Number of ACL filters defined on CPM
=====
Owner MAC IP IPv6 Total

Configuration 0 7 0 7
Host Common 0 2 0 2
Openflow 0 2 1 3

Total 0 14 4 18
=====
Available filters (except openflow): 16369
Available openflow filters: 16381

=====
Number of ACL filter entries / subentries defined on CPM
=====
Inserted by MAC IP IPv6 Total

User configuration 0 21 1 22
 0 21 1 22
Radius 0 0 0 0
 0 0 0 0
Credit Control 0 0 0 0
 0 0 0 0
Embedded 0 0 0 0
 0 0 0 0
Radius shared host 0 2 0 2
 0 2 0 2
Openflow 0 0 0 0
 0 0 0 0
PCC-Rule 0 0 0 0
 0 0 0 0
```

```

Other 0 0 0 0
 0 0 0 0

Total 0 23 1 24
 0 23 1 24
=====
Available subentries (except openflow): 262120
Available openflow subentries: 262144

=====
Filters utilizing most resources (ordered by CPM entries)
=====
Type Id Entries Subentries TCAM entries
 (per FlexPath)

No Mac filters found

Ip 100 5 5 5
Ip 65535 5 5 5
Ip 1 4 4 4
Ip 5:23 2 2 2
Ip 6:24 2 2 2

Ipv6 fSpec-0 0 0 0
Ipv6 fSpec-2345 0 0 0
Ipv6 _tmnx_ofs_system:1 0 0 0
No more Ipv6 filters
=====

=====
Filters utilizing most resources (ordered by CPM subentries)
=====
Type Id Entries Subentries TCAM entries
 (per FlexPath)

No Mac filters found

Ip 100 5 5 5
Ip 65535 5 5 5
Ip 1 4 4 4
Ip 5:23 2 2 2
Ip 6:24 2 2 2

Ipv6 fSpec-0 0 0 0
Ipv6 fSpec-2345 0 0 0
Ipv6 _tmnx_ofs_system:1 0 0 0
No more Ipv6 filters
=====

```

## dest-tracking

**Syntax** **dest-tracking {sap | sdp | ip | ipv6} [detail]**

**Context** tools>dump>filter>resources

- Description** This command displays information about resources pertaining to tracked targets.
- Parameters** **sap | sdp| ip | ipv6** — displays information about SAP, SDP, IPv4, or IPv6 targets  
**detail** — displays detailed information
- Output** The following output is an example of filter resource SAP destination tracking information.

**Sample Output**

```
dest-tracking sap
=====
Unique SAPs with tracked forwarding states =====
Used : 1
Free : 4095
Total : 4096
=====
```

The following output is an example of filter resource SAP destination tracking detailed information.

```
dest-tracking# sap detail
=====
Unique SAPs with tracked forwarding states =====
Num Destination Ref. count

 1 sap 1/2/2 1
=====
Unique SAPs with tracked forwarding states =====
Used : 1
Free : 4095
Total : 4096
=====
```

**egress-pbr**

- Syntax** **egress-pbr [detail]**
- Context** tools>dump>filter>resources
- Description** This command displays the number of allocated unique egress PBR destinations.
- Parameters** **detail** — displays number of allocated unique egress PBR destinations together with a list of destinations and their ref counts
- Output** The following output is an example of filter resource egress PBR destination information.

**Sample Output**

```
*A:Dut-C>tools dump filter resources egress-pbr
=====
Egress PBR destinations
```



```

=====
Name Count

All destinations 8
Unique destinations 4
=====

*A:Dut-C# tools dump filter resources egress-pbr detail

=====
Unique egress PBR destinations
=====
Num Action Ref. count Parameters

 1 Esi L3 1 esi 00:00:00:00:00:00:00:00:01
 ip 5.5.1.5
 if VasToFromAccess
 rtr 123

 2 Esi L3 2 esi 00:00:00:00:00:00:00:00:02
 ip 5.5.0.5
 if VasToFromNetwork
 rtr 123

 3 Red-pol 3 name egress-pbr

 4 Red-pol 2 name ingress-pbr
=====

```

## http-redirect

|                    |                                                                                                                                                                     |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>http-redirect [detail]</b>                                                                                                                                       |
| <b>Context</b>     | tools>dump>filter>resources                                                                                                                                         |
| <b>Description</b> | This command displays the number of unique and total installed HTTP redirect destinations per system.                                                               |
| <b>Parameters</b>  | <b>detail</b> — Displays the number of unique and total installed HTTP redirect destinations per system, together with a list of unique HTTP redirect destinations. |
| <b>Output</b>      | The following output is an example of filter resource HTTP redirect information.                                                                                    |

### Sample Output

```

A:SROS# /tools dump filter resources http-redirect detail
=====
Unique http-redirects
=====
Num Not emb. Emb. Override URL

 1 1 0 no ovr. http://portal_1.com

```

```

2 10 1 no ovr. http://portal_2.com

3 1 0 no ovr. http://portal_3.com

4 1 0 no ovr. http://portal_4.com

5 1 0 no ovr. http://portal_5.com
=====
Http-redirects
=====
Name Count

Unique redirects 5
Not embedded redirects 14
Embedded redirects 1
All redirects 15
=====

```

## iom

|                    |                                                                                                                                                                                                                   |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>iom</b> [ <i>slot-number</i> ]                                                                                                                                                                                 |
| <b>Context</b>     | tools>dump>filter>resources                                                                                                                                                                                       |
| <b>Description</b> | This command shows information about filter resource utilization on all IOMs or a specified IOM. Resource utilization per filter type is available, as well as filters using most resources on a given line card. |
| <b>Parameters</b>  | <i>slot-number</i> — specifies that only the filter resource utilization associated with the IOM card in this slot will be displayed                                                                              |
| <b>Values</b>      | 1 to 10                                                                                                                                                                                                           |
| <b>Output</b>      | The following output is an example of filter resource utilization information for all IOMs.                                                                                                                       |

### Sample Output

```

*A:Dut-C>tools>dump>filter>resources># iom

=====
Number of ACL filter entries used / available on IOMs
=====
Slot Used Available

1 11 65524
2 5 65530
3 5 65530
=====

=====
Number of ACL filters and filter entries used / available on FlexPaths
=====
Slot FlexPath Dir Filters Filters MAC/IP MAC/IP IPv6 IPv6
 used avail entries entries entries entries

```

```

 used avail used avail

1 1 Ingr 2 2045 10 65526 2 28670
 Egr 2 2045 5 32763 2 16382
2 1 Ingr 4 2043 7 65529 2 28670
 Egr 0 2047 2 32766 2 16382
3 1 Ingr 0 2047 7 65529 2 28670
 Egr 0 2047 2 32766 2 16382
=====

Filters utilizing most resources (ordered by TCAM entries per FlexPath)
Only filters present on any IOM are displayed
=====
Type Id Entries Subentries TCAM entries
 (per FlexPath)

No Mac filters found

Ip 100 5 5 5
Ip 5:23 2 2 2
Ip 6:24 2 2 2
Ip 3 1 1 1
Ip 4 1 1 1

Ipv6 fSpec-0 0 0 0
Ipv6 fSpec-2345 0 0 0
No more Ipv6 filters
=====

```

## ip

|                    |                                                                                                                                                                                                |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>ip &lt;filter-id&gt;</b>                                                                                                                                                                    |
| <b>Context</b>     | tools>dump>filter>resources                                                                                                                                                                    |
| <b>Description</b> | This command displays information about the specified IP filter including resource utilization on CPM and IOM, the IOMs on which the filter is used, and the entries using the most resources. |
| <b>Parameters</b>  | <b>filter-id</b> — specifies that only the filter resource utilization associated with this IP filter will be displayed<br><b>Values</b> 1 to 65535                                            |
| <b>Output</b>      | The following output is an example of IP filter resource utilization information.                                                                                                              |

### Sample Output

```

*A:Dut-C>tools>dump>filter>resources># ip 100
=====
Resource utilization details for Ip filter 100
=====

```

```

CPM entries used : 5
CPM subentries used : 5
TCAM entries used (per FlexPath) : 5
Associated with IOMs : 1,2,3,4,5,6,7,8,9,10

```

```

Largest 5 entries

```

| Entry ID | Active | TCAM entries<br>(per FlexPath) |
|----------|--------|--------------------------------|
| 3        | Yes    | 1                              |
| 4        | Yes    | 1                              |
| 5        | Yes    | 1                              |
| 6        | Yes    | 1                              |
| 100      | Yes    | 1                              |

```

=====

```

## ipv6

- Syntax** `ipv6 <filter-id>`
- Context** `tools>dump>filter>resources`
- Description** This command displays information about the specified IPv6 filter including resource utilization on CPM and IOM, the IOMs on which the filter is used, and the entries using the most resources.
- Parameters** **filter-id** — specifies that only the filter resource utilization associated with this IPv6 filter will be displayed
- Values** 1 to 65535
- Output** The following output is an example of IPv6 filter resource utilization information.

### Sample Output

```
*A:Dut-C>tools>dump>filter>resources># ipv6 "fSpec-0"
```

```

=====
Resource utilization details for Ipv6 filter fSpec-0
=====

```

```

CPM entries used : 0
CPM subentries used : 0
TCAM entries used (per FlexPath) : 0
Associated with IOMs : 2

```

```

Largest 5 entries

```

| Entry ID | Active | TCAM entries<br>(per FlexPath) |
|----------|--------|--------------------------------|
|----------|--------|--------------------------------|

```
No Matching Entries

=====
```

## mac

|                    |                                                                                                                                                                                                 |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>mac</b> <filter-id>                                                                                                                                                                          |
| <b>Context</b>     | tools>dump>filter>resources                                                                                                                                                                     |
| <b>Description</b> | This command displays information about the specified MAC filter including resource utilization on CPM and IOM, the IOMs on which the filter is used, and the entries using the most resources. |
| <b>Parameters</b>  | <b>filter-id</b> — specifies that only the filter resource utilization associated with this IPv6 filter will be displayed                                                                       |
| <b>Values</b>      | 1 to 65535                                                                                                                                                                                      |
| <b>Output</b>      | The following output is an example of MAC filter resource utilization information.                                                                                                              |

### Sample Output

```
*A:Dut-C>tools>dump>filter>resources># mac 1

=====
Resource utilization details for Mac filter 1
=====
CPM entries used : 1
CPM subentries used : 1
TCAM entries used (per FlexPath) : 1
Associated with IOMs : 1

Largest 5 entries

Entry ID Active TCAM entries
 (per FlexPath)

1 Yes 1
No more entries defined

=====
```

## sticky-dest

|                    |                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>sticky-dest</b>                                                                          |
| <b>Context</b>     | tools>dump>filter>resources                                                                 |
| <b>Description</b> | This command displays information about resources pertaining to sticky destinations timers. |

**Output** The following output is an example of sticky destination filter action information.

**Sample Output**

```
=====
Filter action - Sticky-dest resources
=====
Used : 1
Free : 2047
Total : 2048
=====
```

## activate-best-dest

**Syntax** **activate-best-dest**

**Context** tools>perform>filter>redirect-policy

**Description** This command allows the operator to force a PBR switch to the best destination selected by the redirect policy when that destination is not currently active as result of sticky destination functionality being enabled for the specified redirect policy. If **hold-time-up** is running, the timer is also expired.

## activate-primary-action

**Syntax** **activate-primary-action**

**Context** tools>perform>filter>ip-filter>entry  
tools>perform>filter>ipv6-filter>entry  
tools>perform>filter>mac-filter>entry

**Description** This command allows an operator to activate the primary action for the given filter policy entry. If the primary action is already active, the command has no effect. If a secondary action is active, the primary action will be activated unless the primary target is down. If the sticky destination timer is running for the primary action entry, it will expire.

## **5 Hybrid OpenFlow Switch**

### **5.1 In This Chapter**

Nokia supports Hybrid OpenFlow Switch (H-OFS) functionality. The hybrid model allows operators to deploy Software Defined Networking (SDN) traffic steering using OpenFlow (OF) atop of the existing routing/switching infrastructure.

## 5.2 Hybrid OpenFlow Switching

The hybrid OpenFlow model allows operators to deploy SDN traffic steering using OpenFlow on top of the existing routing and switching infrastructure. Some of the main benefits of the hybrid model include:

- Increased flexibility and speed for new service deployment—H-OFS implements flexible, policy-driven, standard-based H-OFS traffic steering that allows deployment of new services and on-demand services through policy updates rather than service and infrastructure programming.
- Evolutionary CAPEX/OPEX-optimized SDN deployment—The H-OFS functionality can be deployed on the existing hardware through software upgrade to realize the benefits of FlexPath programmability. The OpenFlow traffic placement is focused access only (that is, flexible, fast, on-demand service deployment) while network infrastructure provides robustness, resiliency, scale, and security.

In a basic mode of operation, a single OpenFlow Switch instance is configured on the router and controlled by a single OpenFlow controller.

The OF controllers and router exchange OpenFlow messages using the OpenFlow protocol (version 1.3.1) over the TCP/IPv4 control channel. Both out-of-band (default) and in-band management is supported for connectivity to the controller. An OpenFlow message is processed by the OpenFlow switch instance on the router that installs all supported H-OFS traffic steering rules in a flow table for the H-OFS instance. A single table per H-OFS instance is supported.

The H-OFS allows operators to:

- Steer IPv4/IPv6 unicast traffic arriving on a Layer 3 interface by programming the 7450 ESS, 7750 SR, and 7950 XRS L3 PBR ACL actions.
- Steer IPv4/IPv6 unicast traffic arriving on a Layer 2 interface by programming the 7450 ESS, 7750 SR, and 7950 XRS L2 PBF ACL actions.
- Drop traffic by programming ACL action drop.
- Forward traffic using regular processing by programming ACL action forward.

Steering actions programmed using OpenFlow are functionally equivalent to ACL actions.

The router allows operators to control traffic using OF, as follows:

- An operator can select a subset of interfaces on the router to have OF rules enabled, by embedding a specific instance of H-OFS in filter policies used only by those interfaces.



- For the interfaces with an H-OFS instance enabled, an operator can:
  - Steer all traffic arriving on an interface by programming the flow table with a “match all” entry.
  - Steer a subset of traffic arriving on an interface with this H-OFS instance enabled by programming the flow table with match rules that select a subset of traffic (OpenFlow match criteria are translated to ACL filter match criteria). Unless explicitly listed as a limitation, the SR OS H-OFS supports any OpenFlow match criteria that can be translated to ACL IPv4/IPv6 filter policy match criteria. A default rule can be assigned for packets that do not match specific rules. These packets can be dropped, forwarded, or sent to the OpenFlow controller.

To enable rules in an H-OFS on an existing service router interface, an operator must:

1. Create one or more ingress line card policies.
2. Assign those line card ingress filter policies to the 7450 ESS, 7750 SR, and 7950 XRS service router interfaces.
3. Embed an H-OFS instance into those line card policies.
4. Program OF rules as required.

OpenFlow can be embedded in IPv4/IPv6 ACL filter policies deployed on:


- L3 IES service interfaces
- L3 network interfaces in base router context
- L3 VPRN service interfaces, including those with NAT
- L2 VPLS service interfaces
- IES/VPRN r-VPLS service interfaces, including those with NAT
- System ACL filters

OpenFlow functionality can be enabled with no effect on forwarding performance. Operators can move from CLI/SNMP programmed steering rules to OpenFlow operational model in service without service disruption.

## 5.2.1 Redundant Controllers and Multiple Switch Instances

The operator can configure one or more instances of an H-OFS (using SNMP or CLI interfaces) with each instance controlled by an OF controller over a unique OF channel using OpenFlow protocol. One OF controller can control multiple H-OFS instances using dedicated channels, or a dedicated OF controller can be deployed per switch. For each switch, up to two OF controllers can be deployed for redundancy. If two controllers are programmed, they can operate in either OFPCR\_ROLE\_EQUAL roles or in OFPCR\_ROLE\_MASTER and OFPCR\_ROLE\_SLAVE roles. [Figure 37](#) shows this architecture.

**Figure 37** SR OS/Switch OF Controller/Switch Architecture Overview



## 5.2.2 GRT-only and Multi-Service H-OFS Modes of Operations

SR OS supports two modes of operation for an H-OFS instance: GRT-only and multi-service. The modes of operation are operator-controlled per H-OFS instance by enabling or disabling the **switch-defined-cookie** option (**configure>open-flow>of-switch>flowtable 0**). For backward compatibility, GRT-only mode of operation is default but, because multi-service mode is a functional superset, Nokia recommends operating in multi-service mode whenever possible. The operator can change the mode in which an H-OFS instance operates but a shutdown is required first. This will purge all the rules forcing the OF controller to reprogram the switch instance after it is re-enabled in a new mode. SR OS supports both H-OFS modes of operation concurrently for different switch instances.

Multi-service modes of operation uses part of the FlowTable cookie field (higher-order 32 bits) to provide the enhanced functionality; the lower-order FlowTable cookie bits are fully controlled by the OF controller. [Table 72](#) depicts higher-order bit Flow Table cookie encoding used when operating in the multi-service mode of operation.

**Table 72 Multi-Service Mode — Higher-Order Bit Flow Table Cookie Encoding**

| sros-cookie Name | sros-cookie Type (Bits 63...60) | sros-cookie Value (Bits 59...32)             | FlowTable Entry Interpretation Based on the sros-cookie                                             |
|------------------|---------------------------------|----------------------------------------------|-----------------------------------------------------------------------------------------------------|
| grt              | 0000                            | 0                                            | FlowTable rule is applicable to GRT instance (IES and router interfaces)                            |
| system           | 1000                            | 0                                            | FlowTable rule is applicable to system filters                                                      |
| service          | 1100                            | service-id for existing VPLS or VPRN service | FlowTable rule is applicable to an existing VPRN or VPLS service specified by the sros-cookie value |

To enable multi-service mode of operation, an operator must embed the OF switch in an ACL filter policy, and, because multi-service H-OFS supports a mix of VPRN/VPLS/GRT/System rules, an additional scope of embedding must be selected (embed open-flow service, embed open-flow system - grt scope used by default). After embedding H-OFS instance, an ACL policy contains rules specific to a VPRN or VPLS service instance or to a GRT or to a System Filter Policy. Therefore, the ACL filter policy can only be used in the scope defined by H-OFS embedding.

Rules programmed by an OF controller with grt, system, and service cookies specified are accepted even if the H-OFS instance is not embedded by a filter activated in a specific context. Rules programmed by an OF controller with a service cookie specified, when the service ID is not one of the supported service types, or when the service with the specified ID does not exist, are rejected with an error returned back to the controller. If an H-OFS is embedded into a line card policy with a specific service context, the embedding must be removed before that service is deleted.

Table 73 summarizes the main differences between the two modes of operation.

**Table 73 Differences Between GRT Mode and Multi-service Mode**

| Function                                                             | GRT Mode<br>(no switch-defined-cookie)                   | Multi-service Mode<br>(switch-defined-cookie)                                                         |
|----------------------------------------------------------------------|----------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| Support OF on IES access interfaces                                  | Yes                                                      | Yes                                                                                                   |
| Support OF on router interfaces in GRT instance                      | Yes                                                      | Yes                                                                                                   |
| Support OF on VPRN access and network interfaces                     | No (lack of native OF service virtualization)            | Yes                                                                                                   |
| Support OF on VPLS access and network interfaces                     | No (lack of native OF service virtualization)            | Yes                                                                                                   |
| Support port and VLAN in flowtable match (see the following section) | No                                                       | Yes                                                                                                   |
| Support OF control of System ACL policies                            | No                                                       | Yes                                                                                                   |
| Traffic steering actions                                             | Forward, drop, redirect to LSP, Layer 3 PBR actions only | All                                                                                                   |
| Scale                                                                | Up to ingress ACL filter policy entry scale              | Up to OF system scale limit per H-OFS instance, and up to 64 534 entries per unique sros-cookie value |

Restrictions:

- Refer to the SR OS R15.0.Rx Software Release Notes for a full list of GRT/IES/VPRN/VPLS interfaces that support OF control for multi-service mode.
- The 7450 ESS, 7750 SR, and 7950 XRS H-OFS always requires sros-cookie to be provided for FlowTable operations and will fail any operation without the cookie when the **switch-defined-cookie** command is enabled.

- OF no-match-action is not programmed in hardware for system filters, because system filters are chained to other filter policies and no-match-action would break the chaining.
- An H-OFS instance does not support overlapping of priorities (flow\_priority value) within a single sros-cookie (type+value). The supported values for priority differ based on a value for **switch-define-cookie**:
  - H-OFS with the **switch-defined-cookie** command disabled
    - Valid flow\_priority\_range 1 to max-size – 1
    - flow\_priority\_value 0 is reserved (no match action)
  - H-OFS with the **switch-defined-cookie** command enabled
    - Valid flow\_priority\_range 1 to 65534
    - flow\_priority\_value 0 is reserved (no match action)
- flow\_priority must map to a valid filter ID. The following items show how flow\_priority is mapped to a filter policy entry ID:
  - H-OFS with **switch-define-cookie** disabled
    - filter entry ID = max-size – flow\_priority + embedding offset
  - H-OFS with **switch-define-cookie** enabled
    - filter entry ID = 65535 – flow\_priority + embedding offset
- When multiple H-OFS instances are embedded into a single ACL filter, no two H-OFS instances can program the same filter entry ID.

### 5.2.2.1 Port and VLAN ID Match in Flow Table Entries

When operating in multi-service mode, SR OS H-OFS supports matching on port and VLAN IDs as part of Flow Table match criteria. When an OF controller specifies incoming port and VLAN values other than "ANY", the H-OFS instance translates them to an SR OS VPLS SAP (sros-cookie must be set to a valid VPLS service ID). If the translation does not result in an existing VPLS SAP, the rule is rejected and an error is returned to the controller.

A flow table rule with a port/VLAN ID match is programmed only if the matching SAP has this H-OFS instance embedded in its ACL ingress filter policy using SAP scope of embedding (**embed open-flow sap**). See [SR OS H-OFS Port and VLAN Encoding](#) for required encoding of port and VLAN IDs.

The SR OS H-OFS supports a mix of rules with service scope and with SAP scope. For VPLS SAPs, an H-OFS instance must be embedded twice: once for the VPLS service and once for the SAP if both service-level and SAP-level rules are to be activated.

An example of activating both service-level and SAP-level rules inside a single ACL policy 1 used on VPLS SAP 1/1/1:100 is as follows:

```
configure filter ip-filter 1
 scope exclusive
 embed open-flow "ofs1" service vpls100 offset 100
 embed open-flow "ofs1" sap 1/1/1:100 offset 200
```

Restrictions:

- Because an H-OFS instance does not support overlapping priorities within a single sros-cookie (type+value), the priority for rules applicable to different SAPs within the same VPLS service must not overlap.
- Masking is not supported when adding a new flow table rule with a port and VLAN ID match.

## 5.2.3 Hybrid OpenFlow Switch Steering using Filter Policies

A router H-OFS instance is embedded into line card IPv4 and IPv6 filter policies to achieve OF-controlled Policy Based Routing (PBR). When H-OFS instance is created, embedded filters (IP and IPv6) required for that instance are automatically created. The filters are created with names, as follows:

“\_tmnx\_ofs\_<ofs\_name>”, with the same name for IPv4 and IPv6 filters used.

If embedded filters cannot be allocated due to the lack of filter policy instances, the creation of an H-OFS instance will fail. When the H-OFS instance is deleted, the corresponding embedded filters are freed.

The H-OFS can be embedded only in ingress filter policies on line cards/platforms supporting embedded filters and for services supporting H-OFS. Embedding of an H-OFS in filter policies on unsupported services is blocked. Embedding of an H-OFS in filter policies in unsupported direction or on unsupported hardware follows the general filter policy misconfiguration behavior and is not recommended. Unsupported match fields are ignored. Other match criteria may cause a packet to match an entry.

As soon as an H-OFS instance is created, the controller can program OF rules for that instance. For instance, the rules can be created prior to the H-OFS instance embedding into a filter policy or prior to a filter policy with H-OFS instance embedded being assigned to an interface. This allows the operator to either preprogram H-OFS steering rules, or to disable the rules without removing them from a flow table by removing the embedding. An error is returned to the controller if it attempts to program rules not supported by the system. The following lists examples of the errors returned:

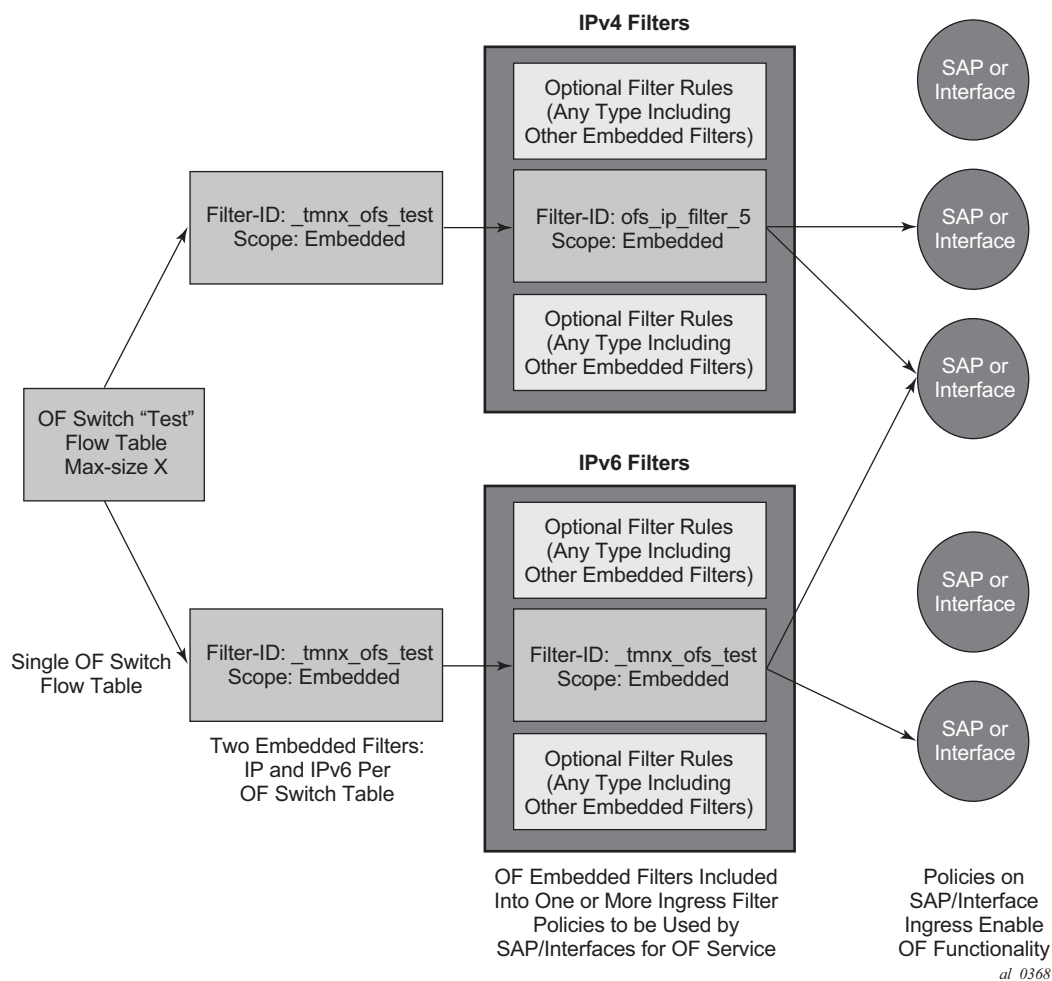
- unsupported instr: [OFPET\_BAD\_INSTRUCTION, OFPBIC\_UNSUP\_INST]
- unsupported action: [OFPET\_BAD\_ACTION, OFPBAC\_BAD\_TYPE]?
- unsupported output port: [OFPET\_BAD\_ACTION, OFPBAC\_BAD\_OUT\_PORT]?
- unsupported match field: [OFPET\_BAD\_MATCH, OFPBMC\_BAD\_FIELD]?
- unsupported match value: [OFPET\_BAD\_MATCH, OFPBMC\_BAD\_VALUE]?
- output port invalid/deleted after flow\_mod is sent to filter: OFPET\_BAD\_ACTION, OFPBAC\_BAD\_OUT\_PORT]?

When the OF controller updates traffic steering rules, the Hybrid OpenFlow Switch updates the flow table rules. This automatically triggers programming of the embedded filter, which consequently causes instantiation of the rules for all services/interfaces that have a filter policy embedding this H-OFS instance. Embedded filter policy configuration/operational rules apply also to embedded filters auto-created for an H-OFS instance (see Embedded Filter Support for ACL Filter Policies section of this guide). MPLS cannot be deleted if OFS rules are created that redirect to an LSP.

The auto-created embedded filters can be viewed through CLI but cannot be modified and/or deleted through filter policy CLI/SNMP. The operator can see the above embedded filters under show filter context, including the details about the filters, entries programmed, interface association, statistics, and so on.

Figure 29 shows the H-OFS to service operator-configurable mapping example.

For an H-OFS with the **switch-defined-cookie** command enabled, embedded filters are created for each unique context in the H-OFS instead.

**Figure 38 OF Flow Table Mapping to Router/Switch Service Infrastructure Example — switch-defined-cookie Disabled**

The router allows mixing H-OFS rules from one or more H-OFS instances in a single filter policy. Co-existence of H-OFS rules in a single policy with CLI/SNMP programmed rules and/or BGP flowspec programmed rules in a single line card filter policy is also supported. When a management interface and an OF controller flow entry have the same filter policy entry, the management interface-created entry overrides the OF controller-created entry; see the embedded filter functional description. For mixing of the rules from multiple management entities, the controller should not program an entry in its Flow Table that would match all traffic, because this would stop evaluation of the filter policy.



The router supports HA for the OF Flow Table content and statistics. On an activity switch, the channel goes down and is reestablished by the newly active CPM. “Fail secure mode” operation takes place during channel reestablishment (OpenFlow rules continue to be applied to the arriving traffic). The OF controller is expected to resynchronize the OF table when the channel is reestablished. On a router reboot or H-OFS instance shutdown, H-OFS Flow Table rules and statistics are purged. An H-OFS instance cannot be deleted unless the H-OFS instance is first removed from all embedding filter policies.

## 5.2.4 Hybrid OpenFlow Switch Statistics

The SR OS Hybrid OpenFlow switch supports statistics retrieval using the OpenFlow protocol. There are two types of statistics that can be collected:

### 1. Statistics for SR OS H-OFS logical ports

Logical port statistics are available for RSVP-TE and MPLS-TP LSP logical ports. The non-zero statistics will be returned as long as an LSP has statistics enabled through an MPLS configuration.

Zero is always returned for logical port statistics for SR-TE LSPs when LSP statistics are not supported on SR-TE LSPs. The statistics can be retrieved regardless of whether an OF switch uses the specified LSP. The returned packet/bytes values are an aggregate of all packets/bytes forwarded over the LSP.

Statistics are not available for any other logical ports encodings.

### 2. Statistics for SR OS H-OFS flow table

Flow table statistics can be retrieved for one or more flow table entries of an H-OFS. The returned packet/bytes values are based on ACL statistics collected in the hardware. An OpenFlow controller can retrieve statistics either directly from hardware or from the ACL CPM-based bulk request cache. The ACL cache is used when processing an OpenFlow statistics multi-part aggregate request message (OFPMMP\_AGGREGATE), or when an OpenFlow statistics multi-part flow message request (OFPM\_FLOW) is translated to multiple flow table entries (a bulk request). When an OpenFlow multi-part flow statistics request message (OFPM\_FLOW) is translated to a single flow table entries request (a single entry request), the counters are read from hardware in real time.

A mix of the two methods can be used to retrieve some flow table statistics from hardware in real time while retrieving other statistics from the cache. See [Filter Policy Statistics](#) for more information about ACL cache and ACL statistics.

When the auxiliary channel is enabled, the switch will set up a dedicated auxiliary channel for statistics. See [OpenFlow Switch Auxiliary Channels](#) for more information.

**Operational Notes:**

- Flow table statistics displayed through the CLI debugging tools (tools>dump>open-flow>of-switch) are read in real time from hardware. However, to protect the system, executing CLI debugging tool commands within 5 s will return the same statistics for any flow that had its statistics read from hardware within the last 5 s.
- When retrieving flow table statistics at scale, Nokia recommends to either use bulk requests, or to pace single entry requests in order to obtain the desired balance between stats real-time accuracy and CPM activity.

## 5.2.5 OpenFlow Switch Auxiliary Channels

The H-OFS supports auxiliary channels, as defined in OpenFlow version 1.3.1. The packet-in and statistics functions are supported on the auxiliary channels as well as on the main channel.

When the auxiliary channel is enabled on a switch (using the **aux-channel-enable** command), the switch will set up a dedicated auxiliary channel for statistics (Auxiliary ID 1) and a dedicated auxiliary channel for packet-in (Auxiliary ID 2) if a packet-in action is configured, to every controller for a given H-OFS switch instance. Auxiliary connections use the same transport as the main connection. The switch will handle any requests over any established channel and respond on the same channel even if a specific requested auxiliary channel is available.

The H-OFS instance uses the packet-in connection for packet-in functionality by default and expects (but does not require) the controller to use the statistics channel for statistics processing by default.

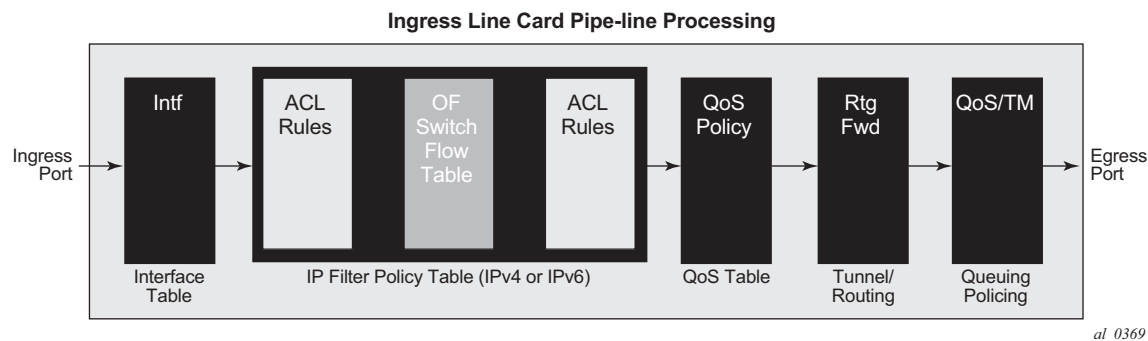
The switch uses the auxiliary channels (packet-in for packet-in-specific requests and statistics for statistics-specific requests) as long as they are available. If they are not available, the switch will use the next available auxiliary channel. If none of the auxiliary channels are available, the main channel will be used.

Auxiliary connections can be enabled or disabled without shutting down the switch.

## 5.2.6 Hybrid OpenFlow Switch Traffic Steering Details

As described in the previous section, an update to an OpenFlow Switch’s flow table results in the embedded filter updates, which triggers an update to all filter policies embedding those filters. The router automatically downloads the new set of rules to the line cards as defined through service configuration. The rules become part of an ingress line card pipeline, as shown in [Figure 39](#).

**Figure 39** OpenFlow Switch Embedding in Ingress Pipeline Processing



### 5.2.6.1 SR OS H-OFS Logical Port

Logical ports are used in OpenFlow to encode switch-specific ports. SR OS H-OFS uses logical ports in steering actions by encoding PBR targets. [Table 74](#) lists logical port types supported by SR OS H-OFS:

**Table 74** Encoding and Supported Logical Port Types

| Bits 31..28             | Bits 27..24                        | Bits 24..0                                                         |
|-------------------------|------------------------------------|--------------------------------------------------------------------|
| Logical port type (LPT) | Logical port type sub-type (LPT-S) | Logical port type value (LPT-V) — always padded with leading zeros |

The following encoding sample shows logical port types supported by SR OS H-OFS:

```
RSVP LSP: LPT: 0100, LPT-S: 0000 (tunnel), LPT-V: RSVP TE Tunnel ID
MPLS-TP LSP: LPT: 0100, LPT-S: 0000 (tunnel), LTP-V: MPLS-TP Tunnel Number
SR-TE LSP: LPT: 0100, LPT-S: 0000 (tunnel), LTP-V: SR-TE LSP-ID
GRT instance: LPT: 0100, LPT-S: 0001 (L3 routing instance), LPT-V: 0
VPRN Id: LPT: 0100, LPT-S: 0001 (L3 routing instance), LPT-V: VPRN Service ID for a
VPRN instance configured on the system, NAT: LPT 0100, LPT-S: 0020 (NAT), LPT-V: 0
```

OF is limited to a 24-bit service ID value range (a subset of VPRN IDs supported by the SR OS system).

Logical port values other than RSVP-TE LSP, SR-TE LSP, and MPLS-TP LSP require H-OFS with the **switch-defined-cookie** command enabled. Only tunnel-encoded ports are stored in the H-OFS logical port table. Therefore, functionality such as retrieving statistics per port is not available for logical ports that are not stored in the H-OFS logical port table.

### 5.2.6.2 SR OS H-OFS Port and VLAN Encoding

The OF controller can use port and VLAN values other than “ANY” for VPLS SAP match and for VPLS steering to SAP for H-OFS instances with the **switch-defined-cookie** command enabled.

To specify a port in an OF message, SR OS TmnxPortId encoding must be used. The allowed values are those for Ethernet physical ports and LAG.

To encode VLAN tags, OXM\_OF\_VLAN\_ID and new experimenter OFL\_OUT\_VLAN\_ID fields are used as shown in [Table 75](#).

**Table 75** VLAN Tag Encoding

|                                            |                                                                            |
|--------------------------------------------|----------------------------------------------------------------------------|
| NULL tag, dot1Q tag, inner QinQ tag VlanId | Outer QinQ tag VlanId                                                      |
| OXM_OF_VLAN_VID                            | OFL_OUT_VLAN_ID (Experimenter field uses same encoding as OXM_OF_VLAN_VID) |

[Table 76](#) shows how OF programmed values are translated to SR OS SAPs.

**Table 76** Translation of OF Programmed Values to SR OS SAPs

| OXM_OF_IN_PORT             | OXM_OF_VLAN_VID               | OFL_OUT_VLAN_ID | Matching SAP SR OS Encoding | Supported in flow_add | Supported in flow_mod flow_del mp_req | Comment             |
|----------------------------|-------------------------------|-----------------|-----------------------------|-----------------------|---------------------------------------|---------------------|
| TmnxPortId for port or LAG | Value: 0x0000<br>Mask: Absent | Must be absent  | port-id<br>lag-id           | Yes                   | Yes                                   | Mask must be absent |

**Table 76 Translation of OF Programmed Values to SR OS SAPs (Continued)**

| OXM_OF_IN_PORT             | OXM_OF_VLAN_VID                                  | OFL_OUT_VLAN_ID                                  | Matching SAP<br>SR OS Encoding                                                           | Supported in<br>flow_add | Supported in<br>flow_mod<br>flow_del<br>mp_req | Comment             |
|----------------------------|--------------------------------------------------|--------------------------------------------------|------------------------------------------------------------------------------------------|--------------------------|------------------------------------------------|---------------------|
| TmnxPortId for port or LAG | Value: 0x1yyy, yyy encodes qtag1<br>Mask: Absent | Must be absent                                   | port-id:qtag1<br>lag-id:qtag1                                                            | Yes                      | Yes                                            | Mask must be absent |
| TmnxPortId for port or LAG | Value: 0x1FFF<br>Mask: Absent                    | Must be absent                                   | port-id:*<br>lag-id:*                                                                    | Yes                      | Yes                                            | Mask must be absent |
| TmnxPortId for port or LAG | Value: 0x1000<br>Mask: 0x1000                    | Must be absent                                   | port-id: any<br>lag-id: any<br>where "any" is either * or a valid VLAN-ID (but not NULL) | No                       | Yes                                            | Mask must be 0x1000 |
| TmnxPortId for port or LAG | Value: 0x1yyy, yyy encodes qtag2<br>Mask: Absent | Value: 0x1zzz, zzz encodes qtag1<br>Mask: Absent | port-id:qtag1.qtag2<br>lag-id:qtag1.qtag2                                                | Yes                      | Yes                                            | Mask must be absent |
| TmnxPortId for port or LAG | Value: 0x1FFF<br>Mask: Absent                    | Value: 0x1zzz, zzz encodes qtag1<br>Mask: Absent | port-id: qtag1.*<br>lag-id: qtag1.*                                                      | Yes                      | Yes                                            | Mask must be absent |
| TmnxPortId for port or LAG | Value: 0x1FFF<br>Mask: Absent                    | Value: 0x1FFF<br>Mask: Absent                    | port-id: *.*<br>lag-id: *.*                                                              | Yes                      | Yes                                            | Mask must be absent |

**Table 76 Translation of OF Programmed Values to SR OS SAPs (Continued)**

| OXM_OF_IN_PORT             | OXM_OF_VLAN_VID               | OFL_OUT_VLAN_ID                                     | Matching SAP SR OS Encoding                                                                        | Supported in flow_add | Supported in flow_mod flow_del mp_req | Comment                                  |
|----------------------------|-------------------------------|-----------------------------------------------------|----------------------------------------------------------------------------------------------------|-----------------------|---------------------------------------|------------------------------------------|
| TmnxPortId for port or LAG | Value: 0x1000<br>Mask: 0x1000 | Value: 0x1zzz,<br>zzz encodes qtag1<br>Mask: Absent | port-id: qtag1.any<br>lag-id: qtag1.any<br>where any is either * or a valid VLAN-ID (but not NULL) | No                    | Yes                                   | Mask must be absent for OFL_OUT_VLAN_VID |
| TmnxPortId for port or LAG | Value: 0x1000<br>Mask: 0x1000 | Value: 0x1FFF<br>Mask: Absent                       | port-id: *.any<br>lag-id: *.any<br>where "any" is either * or a valid VLAN-ID (but not NULL)       | No                    | Yes                                   | Mask must be absent for OFL_OUT_VLAN_VID |
| TmnxPortId for port or LAG | Value: 0x1000<br>Mask: 0x1000 | Value: 0x1000<br>Mask: 0x1000                       | port-id: any.any<br>lag-id: any.any<br>where "any" is either * or a valid VLAN-ID (but not NULL)   | No                    | Yes                                   | Masks must be 0x1000                     |
| TmnxPortId for port or LAG | Value: 0x0000<br>Mask: Absent | Value: 0x1FFF<br>Mask: Absent                       | port-id: *.null                                                                                    | Yes                   | Yes                                   | Mask must be absent                      |

### 5.2.6.3 Redirect to IP next-hop

A router supports redirection of IPv4 or IPv6 next-hop for traffic arriving on a Layer 3 interface. An OF controller can rely on this functionality and program PBR next-hop steering actions for H-OFS instances with the **switch-defined-cookie** command enabled using the following OF encoding:

```

ALU_IPD_EXPERIMENTER_ID: 0x000025BA
ALU_AXN_REDIRECT_TO_NEXTHOP: 2
flow_mod:
instruction= OFPAT_WRITE_ACTION/OFPAT_APPLY_ACTION,
action= OFPAT_EXPERIMENTER(ALU_AXN_REDIRECT_TO_NEXTHOP),
encoding:
struct alu_axn_redirect_to_nhopv4{
uint16_t type; /* OFPAT_EXPERIMENTER. */
uint16_t len; /* Total Length is a multiple of 8. */
uint32_t experimenter; /* Experimenter ID vendor unique*/
uint8_t redirect_type; /* Type = 1 for Nhop*/
uint8_t flags; /* flags is 0-7 bits:
 Bit 0 = Ipv4,
 Bit 1 = Ipv6,
 Bit 2 = indirect
*/
uint8_t pad[2];
uint32_t ipaddr; /* ipv4 addr */
uint8_t pad[0]; /* Not needed */
}; ASSERT(sizeof(alu_axn_redirect_to_nhopv4) == 16)
struct alu_axn_redirect_to_nhopv6{
uint16_t type; /* OFPAT_EXPERIMENTER. */
uint16_t len; /* Total Length is a multiple of 8. */
uint32_t experimenter; /* Experimenter ID vendor unique*/
uint8_t redirect_type; /* Type = 1 for Nhop*/
uint8_t flags; /* flags is 0-7 bits:
 Bit 0 = Ipv4,
 Bit 1 = Ipv6,
 Bit 2 = indirect
*/
uint8_t pad[2];
uint128_t ip6addr; /* ipv6 addr */
uint8_t pad[4]; /* Make total len multiple of 8 */
}; ASSERT(sizeof(alu_axn_redirect_to_nhopv6) == 32)

```

In case of erroneous programming, the following experimenter-specific errors are returned to the controller:

```

enum alu_err_exp_class{
ALU_ERR_CLASS_RD_TO_SDP = 0,
ALU_ERR_CLASS_RD_TO_NHOP = 1,
}
enum alu_err_subtype_redirect_to_nhop
{
ALU_ERR_RN_INVALID_FLAGS = 0
ALU_ERR_RN_INVALID_ARGS = 1
ALU_ERR_RN_INVALID_ADDR = 2
}

```

#### 5.2.6.4 Redirect to GRT Instance or VRF Instance

A router supports redirection of IPv4 or IPv6 traffic arriving on a Layer 3 interface to a different routing instance (GRT or VRF). An OF controller can rely on this functionality and program PBR actions for GRT/VRF steering for H-OFS instances with the **switch-defined-cookie** command enabled using the following OF encoding:

```
flow_mod:
instruction type: OFPAT_WRITE_ACTIONS/OFPAT_APPLY_ACTION,
action type: OFPAT_OUTPUT,
```

port= SR OS LOGICAL port encoding GRT or VPRN Service ID as described in the [SR OS H-OFS Logical Port](#) section.

Because a 24-bit value is used to encode the VPRN service ID in the logical port, redirection to a VPRN service with a service ID above that range is not supported.

#### 5.2.6.5 Redirect to Next-hop and VRF/GRT Instance

A router supports redirection of IPv4 or IPv6 traffic arriving on a Layer 3 interface to a different routing instance (GRT or VRF) and next-hop IP at the same time. An OF controller can rely on this functionality and program PBR steering action for H-OFS instances with the **switch-defined-cookie** command enabled using the following OF encoding:

```
ALU_IPD_EXPERIMENT_ID:0X000025BA
ALU_AXN_REDIRECT_TO_NEXTHOP:2
flow_mod:
Instruction 1:
instruction=OFPAT_WRITE_ACTION/OFPAT_APPLY_ACTION
action=OFPAT_EXPERIMENTER (ALU_AXN_REDIRECT_TO_NEXTHOP) ,
```

Encoding as described in the [Redirect to IP next-hop](#) section (indirect flag must be set).

```
Instruction 2:
instruction type: OFPAT_WRITE_ACTIONS/OFPAT_APPLY_ACTION,
action type: OFPAT_OUTPUT,
```

port= SR OS LOGICAL port encoding GRT or VPRN Service ID as described in the [SR OS H-OFS Logical Port](#) section.



### 5.2.6.6 Redirect to ESI (L2)

The router supports redirection of IPv4/IPv6 traffic arriving on a Layer 2 interface to an Ethernet Segment Identifier (ESI) with an EVPN control plane. An OF controller can program L2 ESI steering with the **switch-defined-cookie** command enabled using the following OF encoding:

```
flow_mod:
 instruction type: OFPAT_WRITE_ACTIONS/OFPAT_APPLY_ACTION,
 action type: OFPAT_EXPERIMENTER(ALU_AXN_REDIRECT_TO_ESI_L2)
 encoding:

struct alu_axn_redirect_to_ESI_L2{
 uint16_t type; /* OFPAT_EXPERIMENTER. */
 uint16_t len; /* Total Length is a multiple of 8. */
 uint32_t experimenter; /* Experimenter ID vendor unique*/
 uint8_t redirect_type ; /* Type = 3 for ESI*/
 uint8_t flags; /* flags is 0-7 bits:
 Value 0 = L2,
 */
 uint8_t esi[10]; /* 10 byte ESI */
 uint32_t svcID; /* Svc-Name Using the OF Encoding */
}; ASSERT(sizeof(alu_axn_redirect_to_ESI_L2) == 24)
```

### 5.2.6.7 Redirect to ESI (L3)

The router supports redirection of IPv4/IPv6 traffic arriving on a Layer 3 interface to an ESI with an EVPN control plane. An OF controller can program L3 ESI steering with the **switch-defined-cookie** command enabled using the following OF encoding:

```
flow_mod:
 instruction type: OFPAT_WRITE_ACTIONS/OFPAT_APPLY_ACTION,
 action type: OFPAT_EXPERIMENTER(ALU_AXN_REDIRECT_TO_ESI_L3)
 encoding:

struct alu_axn_redirect_to_ESI_L3_V4{
 uint16_t type; /* OFPAT_EXPERIMENTER. */
 uint16_t len; /* Total Length is a multiple of 8. */
 uint32_t experimenter; /* Experimenter ID vendor unique*/
 uint8_t redirect_type ; /* Type = 3 for ESI*/
 uint8_t flags; /* flags is 0-7 bits:
 Value 1 = L3 (ipv4)
 */
 uint8_t esi[10]; /* 10 byte ESI */
 uint32_t svcID; /* Svc-Name Using the OF Encoding */
 uint32_t sf-ip; /* v4 address of sf-ip */
 uint32_t ifIndex; /* interface id*/
}; ASSERT(sizeof(alu_axn_redirect_to_ESI_L3_V42) == 32)

struct alu_axn_redirect_to_ESI_L3_V6{
 uint16_t type; /* OFPAT_EXPERIMENTER. */
```

```

uint16_t len; /* Total Length is a multiple of 8. */
uint32_t experimenter; /* Experimenter ID vendor unique*/
uint8_t redirect_type ; /* Type = 1 for Nhop*/
uint8_t flags; /* flags is 0-7 bits:
 Value = 2 = L3 (ipv6)
 */
uint8_t esi[10]; /* 10 byte ESI */
uint32_t svcId; /* Svc-Name Using the OF Encoding */
uint128_t sf-ip; /* v6 address of sf-ip */
uint32_t ifIndex; /* interface id*/
uint8_t pad[4];
}; ASSERT(sizeof(alu_axn_redirect_to_ESI_L3_V6) == 48)

```

### 5.2.6.8 Redirect to ESI IP VAS-Interface Router

The router supports redirection of IPv4/IPv6 traffic arriving on a Layer 3 interface to a VAS interface bound to an ESI with an EVPN control plane. In this encoding, the SF-IP address represents the VAS interface address, and the ifIndex is the VAS interface ID. An OF controller can program L3 steering with the **switch-defined-cookie** command enabled using the following OF encoding:

```

flow_mod:
 instruction type: OFPAT_WRITE_ACTIONS/OFPAT_APPLY_ACTION,
 action type: OFPAT_EXPERIMENTER(ALU_AXN_REDIRECT_TO_ESI_L3)
 encoding:

struct alu_axn_redirect_to_ESI_L3_V4{
 uint16_t type; /* OFPAT_EXPERIMENTER. */
 uint16_t len; /* Total Length is a multiple of 8. */
 uint32_t experimenter; /* Experimenter ID vendor unique*/
 uint8_t redirect_type ; /* Type = 2 for ESI*/
 uint8_t flags; /* flags is 0-7 bits:
 Value 2 = L3 (ipv4)
 */
 uint8_t esi[10];
 uint32_t svcId; /* Svc-Name Using the OF Encoding */
 uint32_t vas-ip; /* v4 address of sf-ip */
 uint32_t ifIndex; /* vas interface id*/
}; ASSERT(sizeof(alu_axn_redirect_to_ESI_L3_V4) == 24)

struct alu_axn_redirect_to_ESI_L3_V6{
 uint16_t type; /* OFPAT_EXPERIMENTER. */
 uint16_t len; /* Total Length is a multiple of 8. */
 uint32_t experimenter; /* Experimenter ID vendor unique*/
 uint8_t redirect_type ; /* Type = 2 for ESI*/
 uint8_t flags; /* flags is 0-7 bits:
 Value 4 = L3 (ipv6)
 */
 uint8_t esi[10]; /* 10 byte ESI */
 uint32_t svcId; /* Svc-Name Using the OF Encoding */
 uint128_t vas-ip; /* v6 address of sf-ip */
 uint32_t ifIndex; /* vas interface id*/
 uint8_t pad[4]
}; ASSERT(sizeof(alu_axn_redirect_to_ESI_L3_V6) == 40)

```

### 5.2.6.9 Redirect to LSP

The router supports traffic steering to an LSP. The following shows the OF encoding to be used by an OF controller:

```
flow_mod:
instruction type: OFFIT_WRITE_ACTIONS or OFFIT_APPLY_ACTION,
action type: OFFPAT_OUTPUT,
```

The port uses SR OS LOGICAL port encoding RSVP-TE, SR-TE, or MPLS-TP LSP as described in the [SR OS H-OFS Logical Port](#) section.

An LSP received in a flow rule is compared against those in the H-OFS logical port table. If the table does not contain the LSP, the rule programming fails. Otherwise, the rule is installed in an ACL filter. As long as any path within the LSP is UP, the redirect rule will forward unicast IPv4 or IPv6 traffic on the current best LSP path by adding an LSP transport label and, in the case of IPv6 traffic, also adding an explicit NULL label.

When an LSP in the H-OFS logical port table goes down, the OF switch removes the LSP from its logical port table and notifies the controller of that fact if the logical port status reporting is enabled. It is up to the OF controller to decide whether to remove rules using this LSP. If the rules are left in the flow table, the traffic that was to be redirected to this LSP will instead be subject to a forward action for this flow rule. If the controller does not remove the entries and the system reuses the LSP identified for another LSP, the rules left in the flow table will start redirecting traffic onto this new LSP.

In some deployments, an SDN controller may need to learn from the router H-OFS logical ports status. To support this function, the OF switch supports optional status reporting using asynchronous OF protocol messages for ports status change.

### 5.2.6.10 Redirect to NAT

The router supports redirection of IPv4 traffic arriving on a Layer 3 interface for ISA NAT processing. An OF controller can program NAT steering for H-OFS instances with the **switch-defined-cookie** command enabled using the following OF encoding:

```
flow_mod:
instruction type: OFFIT_WRITE_ACTIONS/OFFIT_APPLY_ACTION,
action type: OFFPAT_OUTPUT,
```

The port uses SR-OS LOGICAL port encoding as described in the [SR OS H-OFS Logical Port](#) section.

### 5.2.6.11 Redirect to SAP

For traffic arriving on a VPLS interface, a router supports PBF to steer traffic over another VPLS SAP in the same service. An OF controller can rely on this functionality and program PBF steering action for H-OFS instances with the **switch-defined-cookie** command enabled using the following OF encoding:

```
flow_mod:
instruction type: OFPIT_WRITE_ACTIONS or OFPIT_APPLY_ACTION,
Action 1:
action type: OFFPAT_OUTPUT,
```

The port uses encoding as described in the [SR OS H-OFS Port and VLAN Encoding](#) section.

```
Action 2:
action type=OFFPAT_SET_FIELD
```

OXM TLVs encode SAP VLANs as described in the [SR OS H-OFS Port and VLAN Encoding](#) section:

```
- OXM_OF_VLAN_VID
- OFL_OUT_VLAN_ID (optional)
```

### 5.2.6.12 Redirect to SDP

For traffic arriving on a VPLS interface, a router supports PBF to steer traffic over a VPLS SDP in the same service. An OF controller can rely on this functionality and program PBF steering action for H-OFS instances with **switched-defined-cookie** enabled using the following OF encoding:

```
ALU_IPD_EXPERIMENTER_ID: 0x000025BA
ALU_AXN_REDIRECT_TO_SDP: 1
flow_mod:
instruction= OFPIT_WRITE_ACTIONS/OFPIT_APPLY_ACTIONS,
action= OFFPAT_EXPERIMENTER(ALU_AXN_REDIRECT_TO_SDP),
encoding:
struct alu_axn_redirect_to_sdp{
uint16_t type; /* OFFPAT_EXPERIMENTER. */
uint16_t len; /* Total Length is a multiple of 8. */
uint32_t experimenter; /* Experimenter ID vendor unique*/
uint8_t redirect_type; /* Type = 0 for SDP*/
uint8_t flags; /
/* Flags that can be used to denote info(reserved)*/
uint16_t sdp-id; /* Sdp-id*/
uint32_t vcId; /* Vc-id*/
uint8_t pad[0]; /* Not needed */
}; ASSERT(sizeof(alu_axn_redirect_to_sdp) == 16)
```

In case of erroneous programming, the following experimenter-specific errors are returned to the controller:

```
enum alu_err_exp_class
{
 ALU_ERR_CLASS_RD_TO_SDP = 0,
 ALU_ERR_CLASS_RD_TO_NHOP = 1,
}
enum alu_err_redirect_to_sdp
{
 ALU_ERR_RS_INVALID_FLAGS = 0
 ALU_ERR_RS_INVALID_ARGS = 1
 ALU_ERR_RS_INVALID_SDP_ID = 2
 ALU_ERR_RS_INVALID_VC_ID = 3
}
```

### 5.2.6.13 Redirect to a Specific LSP Used by a VPRN Service

The router supports traffic steering within a VPRN, enabling the transport tunnels used by the SDP to be used for specific flows redirected from the system-selected default. This redirection enables large bandwidth flows to be moved to an alternative LSP.

For matching ingress traffic on a VPRN, the **switch-defined-cookie** command must be enabled, with the cookie encoded to match the ingress VPRN's service ID.

Traffic can be redirected to the following:

- the default PE and a different LSP
- a different PE and the default LSP
- the default PE and the default LSP (traffic that may otherwise egress a SAP will take a specified BGP next hop)
- the default PE and a different VRF
- a different PE, the default LSP, and a different prefix

Parameters must be matched in the OF encoding to steer traffic.

```
flow_mod:
instruction type: OFPAT_WRITE_ACTIONS/OFPAT_APPLY_ACTION,
```

```
Action 1:
action type: OFPAT_EXPERIMENTER
ALU_IPD_EXPERIMENTER_ID: 0x000025BA
ExpType= ALU_AXN_REDIRECT_TO_NEXTHOP,
```

```
Action 2:
action type: OFPAT_OUTPUT,
```

port= SR-OS LOGICAL port encoding RSVP-TE, MPLS-TP LSP, or segment routing, as described in [SR OS H-OFS Logical Port](#) section.

Action 3 (optional): to redirect to a different VPRN

```
Action 3:
action type: OFFPAT_EXPERIMENTER
ALU_IPD_EXPERIMENTER_ID: 0x000025BA
ExpType= ALU_AXN_REDIRECT_TO_VPRN,
```

Encoding:

```
struct alu_axn_redirect_to_vprn {
 uint16_t type; /* OFFPAT_EXPERIMENTER => ff ff */
 uint16_t len;
 uint32_t experimenter; /
 * Vendor specific experimenter id => 00 00 25 ba */
 uint8_t exp_axn_type; /* type => 03 */
 uint8_t exp_axn_flags; /* flag => any value is accepted */
 uint8_t pad[2]; /* pad => 00 00 */
 uint32_t vprn; /* vprn svc id */
};ASSERT(sizeof(alu_axn_redirect_to_vprn) == 16)
```

Action 4 (optional): to redirect to a different prefix

```
Action 4:
action type: OFFPAT_SET_FIELD
```

Field is an IP destination address. Subnet masks are not supported in the set\_field instruction.

### 5.2.6.14 Forward Action

An OF controller can program forward action, when a specific flow is to be forwarded using regular router forwarding. This would be a default behavior if the filter-policy embedding this OF switch instance has a default-action forward and no filter policy rule matches the flow. To implement forward action, the following OF encoding is used:

```
flow_mod:
instruction type: OFFPAT_WRITE_ACTIONS or OFFPAT_APPLY_ACTION,
action type: OFFPAT_OUTPUT,
port= NORMAL
```

where NORMAL is an OF reserved value.

### 5.2.6.15 Drop Action

An OF controller can program a drop action, when packets of a specific flow are to be dropped. To implement a drop action, the OF encoding is a wildcard rule with empty action-set.

### 5.2.6.16 Default No-match Action

Packets that do not match any of the flow table entries programmed by the controller are subject to a default action. The default action is configurable in the CLI using the **no-match-action** command. Three possible no-match actions are supported: drop, fall-through (packets are forwarded with regular processing by the router), and packet-in.

The packet-in action causes packets that do not match entries in the flow table, as programmed by the OpenFlow controller, to be extracted and sent to the controller in a flow-controlled manner. Because EQUAL is supported, packet-in messages are sent to all controllers in the UP state. To protect the controller, only the first packet of a specific 5-tuple flow (source IP address, destination IP address, source port, destination port, protocol) to which the no-match action is applied is sent to the controller. This 5-tuple flow context ages out after 10 s. Each switch instance maintains contexts for up to 8192 outstanding packet-in messages to the controller. If the packet-in action is used, an auxiliary channel should be enabled for packet-in processing (using the **aux-channel-enable** command). A count of packets to which packet-in is applied is also available through the OpenFlow channel statistics.

### 5.2.6.17 Programming of DSCP Remark Action

The router supports DSCP remarking of IPv4/IPv6 packets arriving on VPLS, VPRN, GRT, and system interfaces for OFS instances with the **switch-defined-cookie** command enabled using the following OF encoding:

```
flow_mod:
 instruction type: OFPIT_METER
 action type: with the meterId.
```

The meters are configured using meter modification messages, and are configured before the flow messages are sent with meter instruction:

```
typedef struct tOfpMeterModMsg
{
 tOfpMsgHeader msgHdr;
 uint16_t mtrCommand; /* One of OFP_MTR_CMD_*. */
}
```

---

```

 uint16_t mtrConfig; /* bitmap of OFP_MTR_CFG_*. */
 uint32_t mtrId; /* Meter instance. */
 tOfpMeterBandHeader bands[0]; /* The band list length is inferred from
 the length field in the msgHdr. */
 } tOfpMeterModMsg;

typedef struct tOfpMeterBandHeader
{
 uint16_t bandType; /* One of OFP_MTR_BAND_*. */
 uint16_t length; /* Length in bytes of this band. */
 uint32_t rate; /* Rate for this band. */
 uint32_t burstSize; /* Size of bursts. */
} tOfpMeterBandHeader;

typedef enum eOfpMeterBandType
{
 OFP_MTR_BAND_DROP = 1, /* Drop packet. */
 OFP_MTR_BAND_DSCP_REMARK = 2, /* Remark DSCP in the IP header. */
 OFP_MTR_BAND_EXPERIMENTER = 0xFFFF /* Experimenter meter band. */
} eOfpMeterBandType;

typedef struct tOfpMeterBandDscpRemark
{
 tOfpMeterBandHeader bandHdr; /* OFP_MTR_BAND_DSCP_REMARK */
 uint8_t precLevel; /* Number of drop precedence level to add */
 uint8_t pad[3];
} tOfpMeterBandDscpRemark;

```



---

## 5.3 Configuration Notes

The following information describes OF implementation restrictions:

- SR OS Hybrid OpenFlow Switch requires a software upgrade only and can be enabled on any SR OS or switch running IOM-2 (with restrictions) or newer line cards. For full functionality, performance, and future scale, IOM3-XP or newer line cards and CPM4 or newer control cards are recommended.
- Some platforms may not support all OF functionality based on the underlying hardware. For example, if the underlying hardware does not support IPv6, OF IPv6 functionality will not be supported. If the underlying hardware does not support redirect to LSP, redirect action will be ignored.
- Each flow in an OF flow table must have unique priority. Overlap is not supported.
- Timed expiry of the flow entries is not supported.
- The implementation is compliant by design with OpenFlow specification as applicable to supported router functionality only.



## 5.4 OpenFlow Command Reference

- [Command Hierarchies](#)
- [Command Descriptions](#)

### 5.4.1 Command Hierarchies

- [OpenFlow Commands](#)
- [Show Commands](#)
- [Tools Commands](#)

#### 5.4.1.1 OpenFlow Commands

```
config
— open-flow filter-id [create]
 — [no] of-switch ofs-name [ofs-id ofs-id]
 — [no] aux-channel-enable
 — controller ip-address:port
 — no controller
 — description description-string
 — no description
 — echo-interval seconds
 — no echo-interval
 — echo-multiple value
 — no echo-multiple
 — [no] flowtable of-table-id
 — max-size size
 — no max-size
 — no-match-action {drop | fall-through | packet-in}
 — no no-match-action
 — [no] switch-defined-cookie
 — logical-port-status {rsvp-te | mpls-tp | sr-te}
 — no logical-port-status [rsvp-te | mpls-tp | sr-te]
 — [no] shutdown
```

#### 5.4.1.2 Show Commands

```
show
— open-flow
 — of-switch
 — of-switch ofs-name controller ip-address:port detail
```

- **of-switch** *ofs-name* **status controller** [*ip-address:port*]
- **of-switch** *ofs-name* **controller**
- **of-switch** *ofs-name* **flowtable**
- **of-switch** *ofs-name* **status**
- **of-switch** *ofs-name* **port**

### 5.4.1.3 Tools Commands

- tools
  - **dump**
    - **open-flow**
      - **of-switch** *ofs-name* [**flowtable** *of-table-id*] [{**grt** | **system** | **service-id** *service-id*}] [**cookie** *hex-string*] [**priority** *priority*]
      - **of-switch** *ofs-name* [**flowtable** *of-table-id*] **service-id** *service-id* **sap-id** *sap-id* [**cookie** *hex-string*] [**priority** *priority*]
      - **of-switch** *ofs-name* [**flowtable** *of-table-id*] **summary**

## 5.4.2 Command Descriptions

- [Generic Commands](#)
- [Show Commands](#)
- [Debug Commands](#)

### 5.4.2.1 Generic Commands

#### open-flow

|                    |                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>open-flow</b>                                                                                                                                                                   |
| <b>Context</b>     | config                                                                                                                                                                             |
| <b>Description</b> | This command enables configuration content for OpenFlow Hybrid Switch compatibility.<br><br>The <b>no</b> form of the command removes the OpenFlow configuration from the context. |
| <b>Default</b>     | n/a                                                                                                                                                                                |

#### of-switch

**Syntax**    **[no] of-switch** *ofs-name* [**ofs-id** *ofs-id*]

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>open-flow                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <p>This command creates an OpenFlow switch instance.</p> <p>The <b>no</b> form of the command deletes the OpenFlow switch instance from the context.</p>                                                                                                                                                                                                                                                                                                                                                              |
| <b>Default</b>     | no of-switch                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <p><i>ofs-name</i> — specifies the name of the OpenFlow switch instance, a string up to 32 characters</p> <p><i>ofs-id</i> — Specifies the ID of the switch. This is used together with the chassis MAC address to generate the Datapath ID of the OpenFlow switch instance. If it is not configured, then an automatically generated Datapath ID is assigned to the switch according to the order in which the OFS instance came up relative to other OFS instances on the node.</p> <p><b>Values</b> 1 to 65535</p> |

## aux-channel-enable

|                    |                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] aux-channel-enable</b>                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>open-flow>of-switch                                                                                                                                                                                                                                                                                                                                         |
| <b>Description</b> | <p>This command enables auxiliary connections for the given H-OFS instance. If enabled, the H-OFS switch sets up a statistics auxiliary channel (Auxiliary ID 1) and a packet-in auxiliary channel (Auxiliary ID 2) for the main connection to every configured OpenFlow controller.</p> <p>The <b>no</b> form of this command disables auxiliary connections.</p> |
| <b>Default</b>     | no aux-channel-enable                                                                                                                                                                                                                                                                                                                                              |

## controller

|                    |                                                                                                                                                                                                                                                            |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>controller <i>ip-address:port</i></b><br><b>no controller</b>                                                                                                                                                                                           |
| <b>Context</b>     | config>open-flow>of-switch                                                                                                                                                                                                                                 |
| <b>Description</b> | <p>This command configures the OpenFlow controller for this OpenFlow switch. Up to two controllers can be configured per OpenFlow switch instance.</p> <p>The <b>no</b> form of this command deletes the controller for this OpenFlow switch instance.</p> |
| <b>Default</b>     | no controller                                                                                                                                                                                                                                              |
| <b>Parameters</b>  | <i>ip-address:port</i> — Specifies the IP address and TCP port for the OpenFlow channel to the controller.                                                                                                                                                 |

---

## description

|                    |                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>description</b> <i>string</i><br><b>no description</b>                                                                                                                                                                              |
| <b>Context</b>     | config>open-flow>of-switch                                                                                                                                                                                                             |
| <b>Description</b> | <p>This command allows the user to configure a description string for the specified OpenFlow controller instance.</p> <p>The <b>no</b> form of this command deletes the description of the specified OpenFlow controller instance.</p> |
| <b>Default</b>     | no description                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <i>string</i> — specifies a description of the OpenFlow switch instance, a string up to 256 characters                                                                                                                                 |

## echo-interval

|                    |                                                                                                                                                                                                                               |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>echo-interval</b> <i>seconds</i><br><b>no echo-interval</b>                                                                                                                                                                |
| <b>Context</b>     | config>open-flow>of-switch                                                                                                                                                                                                    |
| <b>Description</b> | <p>This command configures the Echo Request interval for monitoring the OpenFlow control channels to the controllers for this OpenFlow switch instance.</p> <p>The <b>no</b> form of this command restores default value.</p> |
| <b>Default</b>     | echo-interval 10                                                                                                                                                                                                              |
| <b>Parameters</b>  | <i>seconds</i> — specifies an interval, in seconds<br><b>Values</b> 1 to 3600                                                                                                                                                 |

## echo-multiple

|                    |                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>echo-multiple</b> <i>value</i><br><b>no echo-multiple</b>                                                                                                                                         |
| <b>Context</b>     | config>open-flow>of-switch                                                                                                                                                                           |
| <b>Description</b> | <p>This command configures the number of consecutive Echo Reply messages that must be lost to declare OF control channel down.</p> <p>The <b>no</b> form of this command restores default value.</p> |
| <b>Default</b>     | echo-multiple 3                                                                                                                                                                                      |

---

|                   |                                                                                                     |
|-------------------|-----------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>value</i> — specifies the threshold value for the number of consecutive Echo Reply messages lost |
| <b>Values</b>     | 3 to 100                                                                                            |

## logical-port-status

|                    |                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>logical-port-status</b> { <b>rsvp-te</b>   <b>mpls-tp</b>   <b>sr-te</b> }<br><b>no logical-port-status</b> [ <b>rsvp-te</b>   <b>mpls-tp</b>   <b>sr-te</b> ]                                                                                                                                                                                                               |
| <b>Context</b>     | config>open-flow>of-switch                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | This command enables status change reporting to the OpenFlow controller for the specified logical port type. To report on multiple logical port types, the command needs to be executed multiple times with different logical port specified as required.<br><br>The <b>no</b> form of this command disables status reporting for specified or all (no argument) logical ports. |
| <b>Default</b>     | no logical-port-status                                                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <b>rsvp-te</b> — enables reporting on RSVP-TE LSP logical ports<br><b>mpls-te</b> — enables reporting on MPLS-TE logical ports<br><b>sr-te</b> — enables reporting on SR-TE logical ports                                                                                                                                                                                       |

## shutdown

|                    |                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] shutdown</b>                                                                                                               |
| <b>Context</b>     | config>open-flow>of-switch                                                                                                         |
| <b>Description</b> | This command administratively enables or disables the OpenFlow switch instance. Disabling the switch purges all flowtable entries. |
| <b>Default</b>     | shutdown                                                                                                                           |

## flowtable

|                    |                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] flowtable</b> <i>of-table-id</i>                                                                                                                                           |
| <b>Context</b>     | config>open-flow>of-switch                                                                                                                                                         |
| <b>Description</b> | This command configures the flow table parameters for this OpenFlow switch instance.<br><br>The <b>no</b> form of this command restores flow table configuration default settings. |
| <b>Default</b>     | no flowtable                                                                                                                                                                       |

---

**Parameters**    *of-table-id* — specifies an identifier of the OpenFlow table, a string up to 256 characters

## max-size

**Syntax**    **max-size** *size*  
**no max-size**

**Context**    config>open-flow>of-switch>flowtable

**Description**    This command configures the size for the specified flow table. The OpenFlow switch instance must be shutdown to modify this parameter.

                  The **no** form of this command restores the default size.

**Default**    1000

**Parameters**    *size* — specifies the maximum size limit for the flow table. The size limit is a total for both IPv4 and IPv6.

**Values**        1 to 262144

**Default**       1000

## no-match-action

**Syntax**    **no-match-action** {**drop** | **fall-through** | **packet-in**}  
**no no-match-action**

**Context**    config>open-flow>of-switch>flowtable

**Description**    This command configures the action for the flow table when a packet does not match any entry for the controller.

                  The **no** form of this command restores the default action.

**Default**    no-match-action fall-through

**Parameters**    **drop** — specifies that packets that do not match entries in the flow table as programmed by the OpenFlow switch will be dropped

**fall-through** — specifies that packets that do not match entries in the flow table as programmed by the OpenFlow switch will be forwarded using regular processing by the router. Fall-through applies if an error occurs that prevents a flow table from being installed in a filter policy.

**packet-in** — specifies that packets that do not match entries in the flow table as programmed by the OpenFlow switch will be extracted and sent to the controller in a flow-controlled manner. If this action is used, an auxiliary channel should be enabled for packet-in processing (using the **aux-channel-enable** command).



## switch-defined-cookie

|                    |                                                                                                                                                                                                    |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>switch-defined-cookie</b><br><b>[no] switch-defined-cookie</b>                                                                                                                                  |
| <b>Context</b>     | config>open-flow>of-switch>flowtable                                                                                                                                                               |
| <b>Description</b> | This command enables OpenFlow switch-defined Flow Table cookie encoding for flowtable 0 that allows multi-service operation.<br><br>The <b>no</b> form of the command disables the above function. |
| <b>Default</b>     | no switch-defined-cookie                                                                                                                                                                           |

### 5.4.2.2 Show Commands

The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

## open-flow

|                    |                                              |
|--------------------|----------------------------------------------|
| <b>Syntax</b>      | <b>open-flow</b>                             |
| <b>Context</b>     | show                                         |
| <b>Description</b> | Displays OpenFlow switch hybrid information. |
| <b>Default</b>     | n/a                                          |

## of-switch

|                    |                                                                                                                                                                                                                                                                                                       |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>of-switch</b><br><b>of-switch ofs-name controller ip-address:port detail</b><br><b>of-switch ofs-name status controller [ip-address:port]</b><br><b>of-switch ofs-name controller</b><br><b>of-switch ofs-name flowtable</b><br><b>of-switch ofs-name status</b><br><b>of-switch ofs-name port</b> |
| <b>Context</b>     | show>open-flow                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | This command displays information related to H-OFS configuration and operations as per the parameters specified.                                                                                                                                                                                      |

If no parameter is specified, this command displays a summary for H-OFS instances configured.

- Default** n/a
- Parameters**
- ofs-name* — specifies the name of the configured H-OFS instance, up to 32 characters
  - controller** *ip-address:port* — displays information on the controller for the specified H-OFS instance
- Values**
- ip-address: a.b.c.d
  - port: 1 to 65535
- detail** — displays detailed information
- status** — displays status information for the specified H-OFS switch or its controller
- flowtable** — displays information about flowtables for the specified H-OFS instance
- port** — displays information about the logical OpenFlow ports registered with the specified H-OFS instance

## Output

### Sample Output

```
*A:Dut-A# show open-flow of-switch "s1" status
=====
Open Flow Switch Information
=====
Switch Name : s1
Data Path ID : 0
Admin Status : Up
Echo Interval : 10 seconds
Echo Multiple : 3
Logical Port Type : all
Buffer Size : 256
Num. of Tables : 1
Description : test-sw1
Capabilities Supp. : flow-stats table-stats port-stats
=====

*A:Dut-A# show open-flow of-switch "s1" controller
=====
Open Flow Controller Summary
=====
IP Address Port

10.20.1.2 6633
10.20.1.3 6633

Number of Controllers : 2

=====

*A:Dut-A# show open-flow of-switch "s1" controller 10.20.1.2:6633 detail
=====
Open Flow Controller Information
=====
IP Address : 10.20.1.2 Port : 6633
Role : equal Generation ID : 0
```

-----  
Open Flow Channel Information  
-----

```

Channel ID : 1 Version : 4
Connection Type : primary Operational Status: Up
Operational Flags : socketStateEstablished helloReceived helloTransmitted
 handshake
Async Fltr Packet In
 (Master or Equal) : tableMiss applyAction
 (Slave) : (Not Specified)
Async Fltr Port Status
 (Master or Equal) : portAdd portDelete portModify
 (Slave) : portAdd portDelete portModify
Async Fltr Flow Rem
 (Master or Equal) : idleTimeOut hardTimeOut flowModDelete groupDelete
 (Slave) : (Not Specified)

Echo Time Expiry : 0d 00:00:10 Hold Time Expiry : 0d 00:00:30
Conn. Uptime : 0d 00:00:00 Conn. Retry : 0d 00:00:00

```

-----  
Open Flow Channel Stats - Channel ID(1)  
-----

| Packet Type     | Transmitted Packets | Received Packets | Error Packets |
|-----------------|---------------------|------------------|---------------|
| Hello           | 1                   | 1                | 0             |
| Error           | 0                   | 0                | 0             |
| Echo Request    | 0                   | 70               | 0             |
| Echo Reply      | 70                  | 0                | 0             |
| Experimenter    | 0                   | 0                | 0             |
| Feat. Request   | 0                   | 1                | 0             |
| Feat. Reply     | 1                   | 0                | 0             |
| Get Cfg Request | 0                   | 1                | 0             |
| Get Cfg Reply   | 1                   | 0                | 0             |
| Set Config      | 0                   | 1                | 0             |
| Packet In       | 0                   | 0                | 0             |
| Flow Removed    | 0                   | 0                | 0             |
| Port Status     | 0                   | 0                | 6             |
| Packet Out      | 0                   | 0                | 0             |
| Flow Modify     | 0                   | 0                | 0             |
| Group Modify    | 0                   | 0                | 0             |
| Port Modify     | 0                   | 0                | 0             |
| Table Modify    | 0                   | 0                | 0             |
| Multipart Req   | 0                   | 0                | 0             |
| Multipart Reply | 0                   | 0                | 0             |
| Barrier Request | 0                   | 0                | 0             |
| Barrier Reply   | 0                   | 0                | 0             |
| Get Q Cfg Req   | 0                   | 0                | 0             |
| Get Q Cfg Reply | 0                   | 0                | 0             |
| Role Request    | 0                   | 0                | 0             |
| Role Reply      | 0                   | 0                | 0             |
| Get Async Req   | 0                   | 0                | 0             |
| Get Async Reply | 0                   | 0                | 0             |
| Set Async       | 0                   | 0                | 0             |
| Meter Modify    | 0                   | 0                | 0             |

```

*A:Dut-A# show open-flow of-switch "s1" flowtable
=====
Flow Table Information
=====
Flow Table ID : 0 Max-Size : 1000
No-Match Action : fall-through Curr Num. of Entries : 1
 Max. Num. of Entries : 54
=====

*A:Dut-A# show open-flow of-switch "s1" port
=====
Open Flow Port Stats
=====
Port ID Port Name Transmitted Packets Transmitted Bytes

1073741825 to_B 0 0
1073741826 to_C 0 0
1073741827 to_D 0 0
1073741828 to_E 0 0
1073741829 to_F 0 0
1073742824 1 0 0
=====

*A:Dut-C# show open-flow of-switch "ofs" controller 1.3.8.8:6633 detail
=====
Open Flow Controller Information
=====
IP Address : 1.3.8.8 Port : 6633
Role : equal
Generation ID : 0

Open Flow Channel Information - Channel ID(2)

Channel ID : 2 Version : 4
Connection Type : primary Operational Status: Up
Auxiliary ID : 0
Source Address : 10.20.1.3 Source Port : 49722
Operational Flags : socket-state-established hello-received hello-transmitted
 handshake
Async Fltr Packet In
 (Master or Equal): table-miss apply-action
 (Slave) : (Not Specified)
Async Fltr Port Status
 (Master or Equal): port-add port-delete port-modify
 (Slave) : port-add port-delete port-modify
Async Fltr Flow Rem
 (Master or Equal): idle-time-out hard-time-out flow-mod-delete group-delete
 (Slave) : (Not Specified)
Echo Time Expiry : 0d 00:00:04 Hold Time Expiry : 0d 00:00:24
Conn. Uptime : 0d 01:27:53 Conn. Retry : 0d 00:00:00

Open Flow Channel Stats - Channel ID(2)

Packet Type Transmitted Packets Received Packets Error Packets

Hello 0 0 0
Error 0 0 0
Echo Request 348 174 0

```

|                 |     |     |   |
|-----------------|-----|-----|---|
| Echo Reply      | 174 | 348 | 0 |
| Experimenter    | 0   | 0   | 0 |
| Feat. Request   | 0   | 0   | 0 |
| Feat. Reply     | 0   | 0   | 0 |
| Get Cfg Request | 0   | 0   | 0 |
| Get Cfg Reply   | 0   | 0   | 0 |
| Set Config      | 0   | 0   | 0 |
| Packet In       | 0   | 0   | 0 |
| Flow Removed    | 0   | 0   | 0 |
| Port Status     | 0   | 0   | 0 |
| Packet Out      | 0   | 0   | 0 |
| Flow Modify     | 0   | 0   | 0 |
| Group Modify    | 0   | 0   | 0 |
| Port Modify     | 0   | 0   | 0 |
| Table Modify    | 0   | 0   | 0 |
| Multipart Req   | 0   | 0   | 0 |
| Multipart Reply | 0   | 0   | 0 |
| Barrier Request | 0   | 0   | 0 |
| Barrier Reply   | 0   | 0   | 0 |
| Get Q Cfg Req   | 0   | 0   | 0 |
| Get Q Cfg Reply | 0   | 0   | 0 |
| Role Request    | 0   | 0   | 0 |
| Role Reply      | 0   | 0   | 0 |
| Get Async Req   | 0   | 0   | 0 |
| Get Async Reply | 0   | 0   | 0 |
| Set Async       | 0   | 0   | 0 |
| Meter Modify    | 0   | 0   | 0 |

## Open Flow Channel Information - Channel ID(3)

|                   |                                                                       |                     |         |
|-------------------|-----------------------------------------------------------------------|---------------------|---------|
| Channel ID        | : 3                                                                   | Version             | : 4     |
| Connection Type   | : auxiliary                                                           | Operational Status: | Up      |
| Auxiliary ID      | : 1                                                                   |                     |         |
| Source Address    | : 10.20.1.3                                                           | Source Port         | : 49748 |
| Operational Flags | : socket-state-established hello-received hello-transmitted handshake |                     |         |

## Async Fltr Packet In

(Master or Equal): table-miss apply-action

(Slave) : (Not Specified)

## Async Fltr Port Status

(Master or Equal): port-add port-delete port-modify

(Slave) : port-add port-delete port-modify

## Async Fltr Flow Rem

(Master or Equal): idle-time-out hard-time-out flow-mod-delete group-delete

(Slave) : (Not Specified)

Echo Time Expiry : 0d 00:00:02 Hold Time Expiry : 0d 00:00:22

Conn. Uptime : 0d 01:27:47 Conn. Retry : 0d 00:00:00

## Open Flow Channel Stats - Channel ID(3)

| Packet Type  | Transmitted Packets | Received Packets | Error Packets |
|--------------|---------------------|------------------|---------------|
| Hello        | 0                   | 0                | 0             |
| Error        | 0                   | 0                | 0             |
| Echo Request | 348                 | 174              | 0             |
| Echo Reply   | 174                 | 348              | 0             |
| Experimenter | 0                   | 0                | 0             |

|                 |   |   |   |
|-----------------|---|---|---|
| Feat. Request   | 0 | 0 | 0 |
| Feat. Reply     | 0 | 0 | 0 |
| Get Cfg Request | 0 | 0 | 0 |
| Get Cfg Reply   | 0 | 0 | 0 |
| Set Config      | 0 | 0 | 0 |
| Packet In       | 0 | 0 | 0 |
| Flow Removed    | 0 | 0 | 0 |
| Port Status     | 0 | 0 | 0 |
| Packet Out      | 0 | 0 | 0 |
| Flow Modify     | 0 | 0 | 0 |
| Group Modify    | 0 | 0 | 0 |
| Port Modify     | 0 | 0 | 0 |
| Table Modify    | 0 | 0 | 0 |
| Multipart Req   | 0 | 0 | 0 |
| Multipart Reply | 0 | 0 | 0 |
| Barrier Request | 0 | 0 | 0 |
| Barrier Reply   | 0 | 0 | 0 |
| Get Q Cfg Req   | 0 | 0 | 0 |
| Get Q Cfg Reply | 0 | 0 | 0 |
| Role Request    | 0 | 0 | 0 |
| Role Reply      | 0 | 0 | 0 |
| Get Async Req   | 0 | 0 | 0 |
| Get Async Reply | 0 | 0 | 0 |
| Set Async       | 0 | 0 | 0 |
| Meter Modify    | 0 | 0 | 0 |

-----  
 Open Flow Channel Information - Channel ID(4)  
 -----

|                        |                                                                             |                     |               |
|------------------------|-----------------------------------------------------------------------------|---------------------|---------------|
| Channel ID             | : 4                                                                         | Version             | : 4           |
| Connection Type        | : auxiliary                                                                 | Operational Status: | Up            |
| Auxiliary ID           | : 2                                                                         |                     |               |
| Source Address         | : 10.20.1.3                                                                 | Source Port         | : 49749       |
| Operational Flags      | : socket-state-established hello-received hello-transmitted handshake       |                     |               |
| Async Fltr Packet In   | (Master or Equal): table-miss apply-action                                  |                     |               |
|                        | (Slave) : (Not Specified)                                                   |                     |               |
| Async Fltr Port Status | (Master or Equal): port-add port-delete port-modify                         |                     |               |
|                        | (Slave) : port-add port-delete port-modify                                  |                     |               |
| Async Fltr Flow Rem    | (Master or Equal): idle-time-out hard-time-out flow-mod-delete group-delete |                     |               |
|                        | (Slave) : (Not Specified)                                                   |                     |               |
| Echo Time Expiry       | : 0d 00:00:01                                                               | Hold Time Expiry    | : 0d 00:00:21 |
| Conn. Uptime           | : 0d 01:27:49                                                               | Conn. Retry         | : 0d 00:00:00 |

-----  
 Open Flow Channel Stats - Channel ID(4)  
 -----

| Packet Type   | Transmitted Packets | Received Packets | Error Packets |
|---------------|---------------------|------------------|---------------|
| Hello         | 0                   | 0                | 0             |
| Error         | 0                   | 0                | 0             |
| Echo Request  | 348                 | 174              | 0             |
| Echo Reply    | 174                 | 348              | 0             |
| Experimenter  | 0                   | 0                | 0             |
| Feat. Request | 0                   | 0                | 0             |
| Feat. Reply   | 0                   | 0                | 0             |

```

Get Cfg Request 0 0 0
Get Cfg Reply 0 0 0
Set Config 0 0 0
Packet In 104420 0 0
Flow Removed 0 0 0
Port Status 0 0 0
Packet Out 0 0 0
Flow Modify 0 0 0
Group Modify 0 0 0
Port Modify 0 0 0
Table Modify 0 0 0
Multipart Req 0 0 0
Multipart Reply 0 0 0
Barrier Request 0 0 0
Barrier Reply 0 0 0
Get Q Cfg Req 0 0 0
Get Q Cfg Reply 0 0 0
Role Request 0 0 0
Role Reply 0 0 0
Get Async Req 0 0 0
Get Async Reply 0 0 0
Set Async 0 0 0
Meter Modify 0 0 0

=====
*A:Dut-C#

```

### 5.4.2.3 Debug Commands

The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

#### open-flow

|                    |                                                            |
|--------------------|------------------------------------------------------------|
| <b>Syntax</b>      | <b>open-flow</b>                                           |
| <b>Context</b>     | tools>dump                                                 |
| <b>Description</b> | This command enables dumping of the open-flow information. |
| <b>Default</b>     | n/a                                                        |

#### of-switch

**Syntax** **of-switch** *ofs-name* [**flowtable** *of-table-id*] [{**grt** | **system** | **service-id** *service-id*}] [**cookie** *hex-string*] [**priority** *priority*]  
**of-switch** *ofs-name* [**flowtable** *of-table-id*] **service-id** *service-id* **sap** *sap-id* [**cookie** *hex-*

*string*] [**priority** *priority*]  
**of-switch** *ofs-name* [**flowtable** *of-table-id*] **summary**

|             |                                                                                                                                                                                                                                                                                                                                 |                                                       |                                                                  |
|-------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------|------------------------------------------------------------------|
| Context     | tools>dump>open-flow                                                                                                                                                                                                                                                                                                            |                                                       |                                                                  |
| Description | This command can be used to dump information for a given open-flow switch or its flowtable. Priority and cookie filters are provided no focus on part of a flow table.                                                                                                                                                          |                                                       |                                                                  |
|             | Usage examples:                                                                                                                                                                                                                                                                                                                 |                                                       |                                                                  |
|             | a. <b>tools&gt;dump&gt;open-flow&gt;of-switch ofs-test</b> — This command displays detailed flow information for a given OpenFlow switch. If the switch has <b>switch-defined-cookie</b> enabled the flows with all cookie-types are displayed.                                                                                 |                                                       |                                                                  |
|             | b. <b>tools&gt;dump&gt;open-flow&gt;of-switch ofs summary</b> — This command displays a summary of each cookie context and the number of flows in it for the switch that has <b>switch-defined-cookie</b> enabled. If <b>switch-defined-cookie</b> is disabled, then the total number of entries is displayed (single context). |                                                       |                                                                  |
|             | c. Options like <b>grt</b> , <b>system</b> , <b>service-id</b> , <b>sap-id</b> , <b>cookie</b> , and <b>priority</b> can be used to limit display entries to the specified options.                                                                                                                                             |                                                       |                                                                  |
| Default     | n/a                                                                                                                                                                                                                                                                                                                             |                                                       |                                                                  |
| Parameters  | <i>ofs-name</i> — specifies the name of the OFS instance, up to 32 characters                                                                                                                                                                                                                                                   |                                                       |                                                                  |
|             | <i>of-table-id</i> — specifies the identifier for the OpenFlow table                                                                                                                                                                                                                                                            |                                                       |                                                                  |
|             | Values                                                                                                                                                                                                                                                                                                                          | 0                                                     |                                                                  |
|             | <i>hex-string</i> — specifies the identifier for the OpenFlow cookies                                                                                                                                                                                                                                                           |                                                       |                                                                  |
|             | Values                                                                                                                                                                                                                                                                                                                          | 0x0 to 0xFFFFFFFFFFFFFFFF                             |                                                                  |
|             | <i>priority</i> — specifies the priority for the OpenFlow switch                                                                                                                                                                                                                                                                |                                                       |                                                                  |
|             | Values                                                                                                                                                                                                                                                                                                                          | 0 to 65535                                            |                                                                  |
|             | <i>service-id</i> — specifies the identifier for the service                                                                                                                                                                                                                                                                    |                                                       |                                                                  |
|             | Values                                                                                                                                                                                                                                                                                                                          | 1 to 2148007978   <i>svc-name</i> : 64 characters max |                                                                  |
|             | <i>sap-id</i> — specifies the identifier for the Ethernet SAP                                                                                                                                                                                                                                                                   |                                                       |                                                                  |
|             | Values                                                                                                                                                                                                                                                                                                                          |                                                       |                                                                  |
|             | <i>sap-id</i>                                                                                                                                                                                                                                                                                                                   | null                                                  | port-id   bundle-id   bpgrp-id   lag-id   aps-id>                |
|             |                                                                                                                                                                                                                                                                                                                                 | dot1q                                                 | port-id   bundle-id   bpgrp-id   lag-id   aps-id   pw-id>:qtag1  |
|             |                                                                                                                                                                                                                                                                                                                                 | qinq                                                  | port-id   bundle-id   bpgrp-id   lag-id   pw-id>:qtag1.qtag2     |
|             |                                                                                                                                                                                                                                                                                                                                 | atm                                                   | <port-id   aps-id>[:vpi/vci   vpi   vpi1.vpi2   cp.conn-prof-id] |
|             |                                                                                                                                                                                                                                                                                                                                 | cp                                                    | keyword                                                          |



|            |                                                          |                    |
|------------|----------------------------------------------------------|--------------------|
|            | conn-prof-id                                             | 1..8000            |
| frame      | port-id   aps-id:dldci                                   |                    |
| cisco-hdlc | slot/mda/port.channel                                    |                    |
| cem        | slot/mda/port.channel                                    |                    |
| ima-grp    | bundle-id>[:vpi/vci   vpi   vpi1.vpi2   cp.conn-prof-id] |                    |
|            | cp                                                       | keyword            |
|            | conn-prof-id                                             | 1..8000            |
| port-id    | slot/mda/port[.channel]                                  |                    |
| bundle-id  | bundle-<type>-slot/mda. <i>bundle-num</i>                |                    |
|            | bundle                                                   | keyword            |
|            | type                                                     | ima, fr, ppp       |
|            | bundle-num                                               | 1..336             |
| bpgrp-id   | bpgrp-<type>-<bpgrp-num>                                 |                    |
|            | bpgrp                                                    | keyword            |
|            | type                                                     | ima, ppp           |
|            | bpgrp-num                                                | 1..2000            |
| aps-id     | aps-<group-id>[.channel]                                 |                    |
|            | aps                                                      | keyword            |
|            | group-id                                                 | 1..64              |
| ccag-id    | <i>ccag-id.path-id[cc-type]&lt;cc-id</i>                 |                    |
|            | ccag                                                     | keyword            |
|            | id                                                       | 1..8               |
|            | path-id                                                  | a, b               |
|            | cc-type                                                  | .sap-net, .net-sap |
|            | cc-id                                                    | 0..4094            |
| eth-tunnel | eth-tunnel- <i>id</i> [:eth-tun-sap-id]                  |                    |
|            | id                                                       | 1..1024            |
|            | eth-tun-sap-id                                           | 0..4094            |
| lag-id     | lag-id                                                   |                    |
|            | lag                                                      | keyword            |
|            | id                                                       | 1..800             |
| pw-id      | <i>pw-id</i>                                             |                    |
|            | pw                                                       | keyword            |
|            | id                                                       | 1..10239           |
| qtag1      | *, 0..4094                                               |                    |
| qtag2      | *   0..4094                                              |                    |
| vpi        | 0..4095 (NNI)                                            |                    |

```

 0..255 (UNI)
vci 1, 2, 5..65535
dlci 16..1022
tunnel-id tunnel-id.private | public:tag
 tunnel
 id
 tag
 keyword
 1..16
 0..4094

```

**summary** — keyword to summarize output

*ip-address:port* — *ip-address*: a.b.c.d

*port*: 1 to 65535

## Output

### Sample Output

```

=====
Switch: ofs
=====
Table : 0 Flow Pri : 0
Cookie : 0x0000000000000000 CookieType: grt
Controller: :::0
Filter Hnd: 0xC30000010000FFFF
Filter : _tmnx_ofs_ofs:1 entry 65535

In Port : *
VID : * Outer VID : *
EthType : *
Src IP : *
Dst IP : *
IP Proto : * DSCP : *
Src Port : * Dst Port : *
ICMP Type : * ICMP Code : *
Label : *
IPv6ExtHdr: (Not Specified)

Action : Fall-through

Flow Flags: IPv4/6 [!E] [RO] [DEF]
Up Time : 0d 00:03:51 Add TS : 680828
Mod TS : 0 Stats TS : 703820
#Packets : 0 #Bytes : 0

Table : 0 Flow Pri : 16
Cookie : 0x0000000000000000 CookieType: grt
Controller: 1.3.8.8:6633
Filter Hnd: 0x830000010000FFEF
Filter : _tmnx_ofs_ofs:1 entry 65519

In Port : *
VID : * Outer VID : *
EthType : 0x0800

```

```

Src IP : *
Dst IP : 22.22.22.1/32
IP Proto : *
Src Port : *
ICMP Type : *
Label : *

DSCP : *
Dst Port : *
ICMP Code : *

Action : Forward On Svc 99

Flow Flags: IPv4
Up Time : 0d 00:01:15
Mod TS : 0
#Packets : 0
Add TS : 696581
Stats TS : 703820
#Bytes : 0

Table : 0
Cookie : 0xC000006300000000
Controller: 1.3.8.8:6633
Filter Hnd: 0x8300000D0000FFEE
Filter : _tmnx_ofs_ofs:13 entry 65518
Flow Pri : 17
CookieType: service 99

In Port : *
VID : *
EthType : 0x0800
Src IP : *
Dst IP : 22.22.22.2/32
IP Proto : *
Src Port : *
ICMP Type : *
Label : *

Outer VID : *
DSCP : *
Dst Port : *
ICMP Code : *

Action : Forward On GRT

Flow Flags: IPv4
Up Time : 0d 00:01:10
Mod TS : 0
#Packets : 0
Add TS : 697095
Stats TS : 703820
#Bytes : 0

Table : 0
Cookie : 0xC00007E200000000
Controller: 1.3.8.8:6633
Filter Hnd: 0x830000050000FFFB
Filter : _tmnx_ofs_ofs:5 entry 65531
SAP : 1/1/3:0
Flow Pri : 4
CookieType: service 2018

In Port : 0x2218000
VID : 0x1000
EthType : 0x0800
Src IP : *
Dst IP : *
IP Proto : *
Src Port : *
ICMP Type : *
Label : *

Outer VID : *
DSCP : *
Dst Port : *
ICMP Code : *

Action : Forward On Sap
 Sap 1/1/3:0

Flow Flags: IPv4
Up Time : 0d 00:02:13
Add TS : 690788

```

```

Mod TS : 0 Stats TS : 703820
#Packets : 0 #Bytes : 0

Table : 0 Flow Pri : 3
Cookie : 0xC00007E200000000 CookieType: service 2018
Controller: 1.3.8.8:6633
Filter Hnd: 0x830000040000FFFC
Filter : _tmnx_ofs_ofs:4 entry 65532
SAP : 1/1/3:4094

In Port : 0x2218000
VID : 0xlffe Outer VID : *
EthType : 0x0800
Src IP : *
Dst IP : *
IP Proto : * DSCP : *
Src Port : * Dst Port : *
ICMP Type : * ICMP Code : *
Label : *

Action : Forward On Sap
 Sap 1/1/3:4094

Flow Flags: IPv4
Up Time : 0d 00:02:18 Add TS : 690274
Mod TS : 0 Stats TS : 703820
#Packets : 0 #Bytes : 0

Table : 0 Flow Pri : 5
Cookie : 0xC00007E200000000 CookieType: service 2018
Controller: 1.3.8.8:6633
Filter Hnd: 0x830000060000FFFA
Filter : _tmnx_ofs_ofs:6 entry 65530
SAP : lag-800:4094

In Port : 0x50000320
VID : 0xlffe Outer VID : *
EthType : 0x0800
Src IP : *
Dst IP : *
IP Proto : * DSCP : *
Src Port : * Dst Port : *
ICMP Type : * ICMP Code : *
Label : *

Action : Forward On Sap
 Sap lag-800:4094

Flow Flags: IPv4
Up Time : 0d 00:02:09 Add TS : 691201
Mod TS : 0 Stats TS : 703821
#Packets : 0 #Bytes : 0

Table : 0 Flow Pri : 8
Cookie : 0xC00007E300000000 CookieType: service 2019
Controller: 1.3.8.8:6633
Filter Hnd: 0x830000090000FFF7
Filter : _tmnx_ofs_ofs:9 entry 65527
SAP : 2/1/3:1.0

```

```

In Port : 0x4218000
VID : 0x1000
EthType : 0x0800
Src IP : *
Dst IP : *
IP Proto : *
Src Port : *
ICMP Type : *
Label : *

Outer VID : 0x1001

DSCP : *
Dst Port : *
ICMP Code : *

Action : Forward On Sap
 Sap 2/1/3:1.0

Flow Flags: IPv4
Up Time : 0d 00:01:56
Mod TS : 0
#Packets : 0
Add TS : 692448
Stats TS : 703821
#Bytes : 0

Table : 0
Cookie : 0xC00007E300000000
Controller: 1.3.8.8:6633
Filter Hnd: 0x830000080000FFF8
Filter : _tmnx_ofs_ofs:8 entry 65528
SAP : 2/1/3:4094.4094

Flow Pri : 7
CookieType: service 2019

In Port : 0x4218000
VID : 0x1ffe
EthType : 0x0800
Src IP : *
Dst IP : *
IP Proto : *
Src Port : *
ICMP Type : *
Label : *

Outer VID : 0x1ffe

DSCP : *
Dst Port : *
ICMP Code : *

Action : Forward On Sap
 Sap 2/1/3:4094.4094

Flow Flags: IPv4
Up Time : 0d 00:02:01
Mod TS : 0
#Packets : 0
Add TS : 692032
Stats TS : 703821
#Bytes : 0

Table : 0
Cookie : 0xC00007E300000000
Controller: 1.3.8.8:6633
Filter Hnd: 0x8300000B0000FFF5
Filter : _tmnx_ofs_ofs:11 entry 65525
SAP : lag-799:4094.4094

Flow Pri : 10
CookieType: service 2019

In Port : 0x5000031f
VID : 0x1ffe
EthType : 0x0800
Src IP : *
Dst IP : *
IP Proto : *
Src Port : *
ICMP Type : *

Outer VID : 0x1ffe

DSCP : *
Dst Port : *
ICMP Code : *

```

```
Label : *

Action : Forward On Sap
 Sap lag-799:4094.4094

Flow Flags: IPv4
Up Time : 0d 00:01:46 Add TS : 693483
Mod TS : 0 Stats TS : 703821
#Packets : 0 #Bytes : 0

Table : 0 Flow Pri : 1
Cookie : 0xC00007E400000000 CookieType: service 2020
Controller: 1.3.8.8:6633
Filter Hnd: 0x830000020000FFFE
Filter : _tmnx_ofs_ofs:2 entry 65534
SAP : 2/1/4

In Port : 0x4220000
VID : 0x0 Outer VID : *
EthType : 0x0800
Src IP : *
Dst IP : *
IP Proto : * DSCP : *
Src Port : * Dst Port : *
ICMP Type : * ICMP Code : *
Label : *

Action : Forward On Sap
 Sap 2/1/4

Flow Flags: IPv4
Up Time : 0d 00:02:27 Add TS : 689443
Mod TS : 0 Stats TS : 703821
#Packets : 0 #Bytes : 0

Table : 0 Flow Pri : 12
Cookie : 0xC00007E400000000 CookieType: service 2020
Controller: 1.3.8.8:6633
Filter Hnd: 0x830000020000FFF3
Filter : _tmnx_ofs_ofs:2 entry 65523
SAP : 2/1/4

In Port : 0x4220000
VID : 0x0 Outer VID : *
EthType : 0x0800
Src IP : *
Dst IP : *
IP Proto : * DSCP : *
Src Port : * Dst Port : *
ICMP Type : * ICMP Code : *
Label : *

Action : Forward Sdp 12:4294967295

Flow Flags: IPv4
Up Time : 0d 00:01:36 Add TS : 694524
Mod TS : 0 Stats TS : 703821
#Packets : 0 #Bytes : 0

```

```

Table : 0
Cookie : 0xC00007E400000000
Controller: 1.3.8.8:6633
Filter Hnd: 0x830000020000FFF2
Filter : _tmnx_ofs_ofs:2 entry 65522
SAP : 2/1/4
Flow Pri : 13
CookieType: service 2020

```

```

In Port : 0x4220000
VID : 0x0
EthType : 0x0800
Src IP : *
Dst IP : *
IP Proto : *
Src Port : *
ICMP Type : *
Label : *
Outer VID : *
DSCP : *
Dst Port : *
ICMP Code : *

```

```

Action : Forward On Nhop(Indirect)
 Nhop: 200.180.200.180

```

```

Flow Flags: IPv4
Up Time : 0d 00:01:31
Mod TS : 0
#Packets : 0
Add TS : 695037
Stats TS : 703821
#Bytes : 0

```

```

Table : 0
Cookie : 0xC00007E400000000
Controller: 1.3.8.8:6633
Filter Hnd: 0x830000020000FFF0
Filter : _tmnx_ofs_ofs:2 entry 65520
SAP : 2/1/4
Flow Pri : 15
CookieType: service 2020

```

```

In Port : 0x4220000
VID : 0x0
EthType : 0x0800
Src IP : *
Dst IP : *
IP Proto : *
Src Port : *
ICMP Type : *
Label : *
Outer VID : *
DSCP : *
Dst Port : *
ICMP Code : *

```

```

Action : Forward LspId 1
 Lsp lsp1

```

```

Flow Flags: IPv4
Up Time : 0d 00:01:21
Mod TS : 0
#Packets : 0
Add TS : 696067
Stats TS : 703822
#Bytes : 0

```

```

Table : 0
Cookie : 0xC00007E400000000
Controller: 1.3.8.8:6633
Filter Hnd: 0x430000020000FFF1
Filter : _tmnx_ofs_ofs:2 entry 65521
SAP : 2/1/4
Flow Pri : 14
CookieType: service 2020

```

```

In Port : 0x4220000
VID : 0x0
Outer VID : *

```

```

EthType : 0x86dd
Src IP : *
Dst IP : *
IP Proto : *
Src Port : *
ICMP Type : *
Label : *
IPv6ExtHdr: (Not Specified)

DSCP : *
Dst Port : *
ICMP Code : *

Action : Forward On Nhop(Indirect)
 Nhop: 3ffe:1111:1111:2222:2222:3333:3333:4444

Flow Flags: IPv6
Up Time : 0d 00:01:26
Mod TS : 0
#Packets : 0
Add TS : 695551
Stats TS : 703822
#Bytes : 0

Table : 0
Cookie : 0xC00007E400000000
Controller: 1.3.8.8:6633
Filter Hnd: 0x830000030000FFFD
Filter : _tmnx_ofs_ofs:3 entry 65533
SAP : lag-798
Flow Pri : 2
CookieType: service 2020

In Port : 0x5000031e
VID : 0x0
EthType : 0x0800
Src IP : *
Dst IP : *
IP Proto : *
Src Port : *
ICMP Type : *
Label : *
Outer VID : *
DSCP : *
Dst Port : *
ICMP Code : *

Action : Forward On Sap
 Sap lag-798

Flow Flags: IPv4
Up Time : 0d 00:02:23
Mod TS : 0
#Packets : 0
Add TS : 689857
Stats TS : 703822
#Bytes : 0

Table : 0
Cookie : 0x8000000000000000
Controller: 1.3.8.8:6633
Filter Hnd: 0x4300000E0000FFEC
Filter : _tmnx_ofs_ofs:14 entry 65516
Flow Pri : 19
CookieType: system

In Port : *
VID : *
EthType : 0x86dd
Src IP : *
Dst IP : 3ffe::1616:1601/128
IP Proto : *
Src Port : *
ICMP Type : *
Label : *
Outer VID : *
DSCP : *
Dst Port : *
ICMP Code : *
IPv6ExtHdr: (Not Specified)

```



```

Action : Forward On Nhop(Indirect)
 Nhop: 3ffe:1111:1111:2222:2222:3333:3333:4444

Flow Flags: IPv6
Up Time : 0d 00:01:01 Add TS : 698121
Mod TS : 0 Stats TS : 703822
#Packets : 0 #Bytes : 0

Table : 0 Flow Pri : 18
Cookie : 0x8000000000000000 CookieType: system
Controller: 1.3.8.8:6633
Filter Hnd: 0x83000000E0000FFED
Filter : _tmnx_ofs_ofs:14 entry 65517

In Port : *
VID : * Outer VID : *
EthType : 0x0800
Src IP : *
Dst IP : 22.22.22.1/32
IP Proto : * DSCP : *
Src Port : * Dst Port : *
ICMP Type : * ICMP Code : *
Label : *

Action : Forward On Nhop(Indirect)
 Nhop: 200.180.200.180

Flow Flags: IPv4
Up Time : 0d 00:01:06 Add TS : 697608
Mod TS : 0 Stats TS : 703822
#Packets : 0 #Bytes : 0

Number of flows: 17
=====
*A:Dut-C#

```



---

## 6 Cflowd

### 6.1 Cflowd Overview

Cflowd is a tool used to sample IPv4, IPv6, MPLS, and Ethernet traffic data flows through a router. Cflowd enables traffic sampling and analysis by ISPs and network engineers to support capacity planning, trends analysis, and characterization of workloads in a network service provider environment.

Cflowd is also useful for traffic engineering, network planning and analysis, network monitoring, developing user profiles, data warehousing and mining, as well as security-related investigations. Collected information can be viewed several ways such as in port, AS, or network matrices, and pure flow structures. The amount of data stored depends on the cflowd configurations.

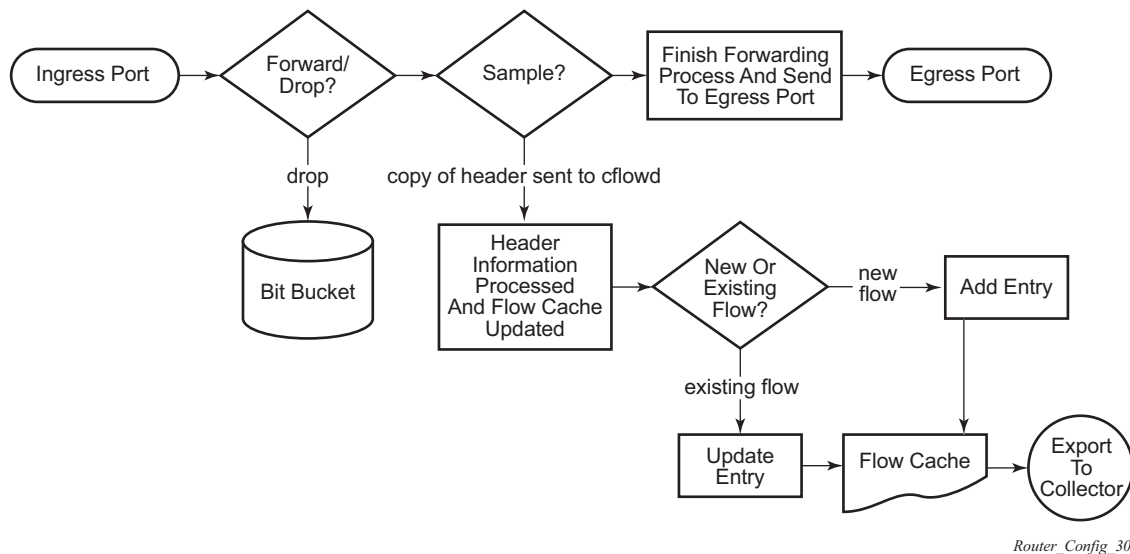
Cflowd maintains a list of data flows through a router. A flow is a unidirectional traffic stream defined by several characteristics such as source and destination IP addresses, source and destination ports, inbound interface, IP protocol and TOS bits.

When a router receives a packet for which it currently does not have a flow entry, a flow structure is initialized to maintain state information regarding that flow, such as the number of bytes exchanged, IP addresses, port numbers, AS numbers, and so on. Each subsequent packet matching the same parameters of the flow contributes to the byte and packet count of the flow until the flow is terminated and exported to a collector for storage.

For the 7450 ESS-7 and 7450 ESS-12, cflowd is only supported if mixed mode is enabled.

#### 6.1.1 Operation

[Figure 40](#) shows the basic operation of the cflowd feature. This sample flow is only used to describe the basic steps that are performed. It is not intended to specify implementation.

**Figure 40 Basic Cflowd Steps**

1. As a packet ingresses a port, a decision is made to forward or drop the packet.
2. If the packet is forwarded, it is then decided whether the packet should be sampled for cflowd.
3. If a new flow is found, a new entry is added to the cache. If the flow already exists in the cache, the flow statistics are updated.
4. If a new flow is detected and the maximum number of entries are already in the flow cache, the earliest expiry entry is removed. The earliest expiry entry/flow is the next flow that will expire due to the active or inactive timer expiration.
5. If a flow has been inactive for a period of time equal to or greater than the inactive timer (default 15 s), the entry is removed from the flow cache.
6. If a flow has been active for a period of time equal to or greater than the active timer (default 30 min), the entry is removed from the flow cache.

When a flow is exported from the cache, the collected data is sent to an external collector, which maintains an accumulation of historical data flows that network operators can use to analyze traffic patterns.

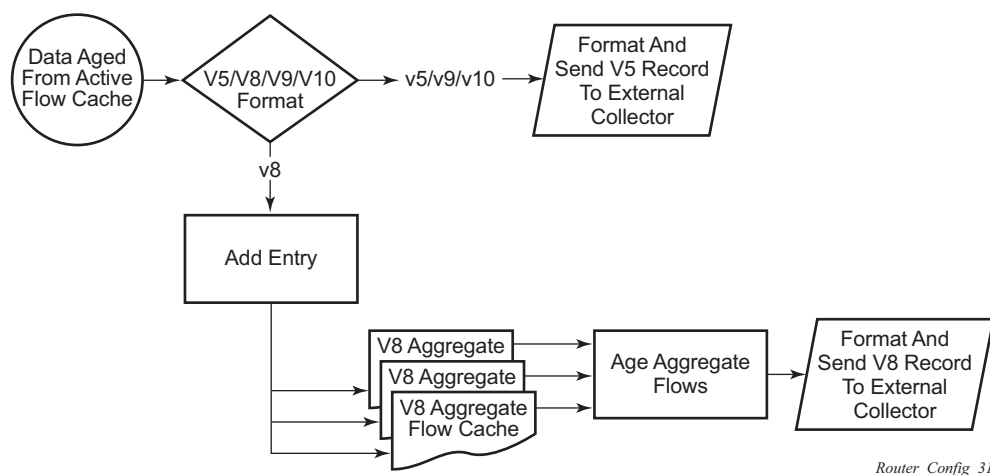
Data is exported in one of the following formats:

- Version 5 — Generates a fixed export record for each individual flow captured.
- Version 8 — Aggregates multiple individual flows into a fixed aggregate record.
- Version 9 — Generates a variable export record, depending on user configuration and sampled traffic type (IPv4, IPv6, or MPLS), for each individual flow captured.

- Version 10 (IPFIX) — Generates a variable export record, depending on user configuration and sampled traffic type (IPv4, IPv6, or MPLS), for each individual flow captured.

Figure 41 shows V5, V8, V9, and V10 flow processing.

**Figure 41 V5, V8, V9, V10, and Flow Processing**



1. As flows are expired from the active flow cache, the export format must be determined, either V5, V8, V9, and V10.
  - If the export format is V5 or V9 and V10, no further processing is performed and the flow data is accumulated to be sent to the external collector.
  - If the export format is V8, the flow entry is added to one or more of the configured aggregation matrices.
  - As the entries within the aggregate matrices are aged out, they are accumulated to be sent to the external flow collector in V8 format.

The sample rate and cache size are configurable values. The cache size default is 64K flow entries.

A flow terminates when one of the following conditions is met:

- When the inactive timeout period expires (default: 15 s). A flow is considered terminated when no packets are seen for the flow for  $n$  seconds.
- When an active timeout expires (default: 30 s). Default active timeout is 30 min. A flow terminates according to the time duration, regardless of whether there are packets coming in for the flow.
- When the user executes a **clear cflowd** command.

- When other measures are met that apply to aggressively age flows as the cache becomes too full (such as overflow percent).

#### **6.1.1.1 Version 8**

There are several different aggregate flow types including:

- AS matrix
- Destination prefix matrix
- Source prefix matrix
- Prefix matrix
- Protocol/port matrix.

Version 8 is an aggregated export format. As individual flows are aged out of the raw flow cache, the data is added to the aggregate flow cache for each configured aggregate type. Each of these aggregate flows are also aged in a manner similar to the method the active flow cache entries are aged. When an aggregate flow is aged out, it is sent to the external collector in the V8 record format.

#### **6.1.1.2 Version 9**

Version 9 format is a more flexible format and allows for different templates or sets of cflowd data to be sent based on the type of traffic being sampled and the template set configured.

Version 9 is interoperable with RFC 3954, *Cisco Systems NetFlow Services Export Version 9*.

#### **6.1.1.3 Version 10**

Version 10 is a new format and protocol that interoperates with the specifications from the IETF as the IP Flow Information Export (IPFIX) standard. Like V9, the V10 format uses templates to allow for different data elements regarding a flow that is to be exported and to handle different type of data flows, such as IPv4, IPv6, and MPLS.

Version 10 is interoperable with RFC 5150 and 5102.

---

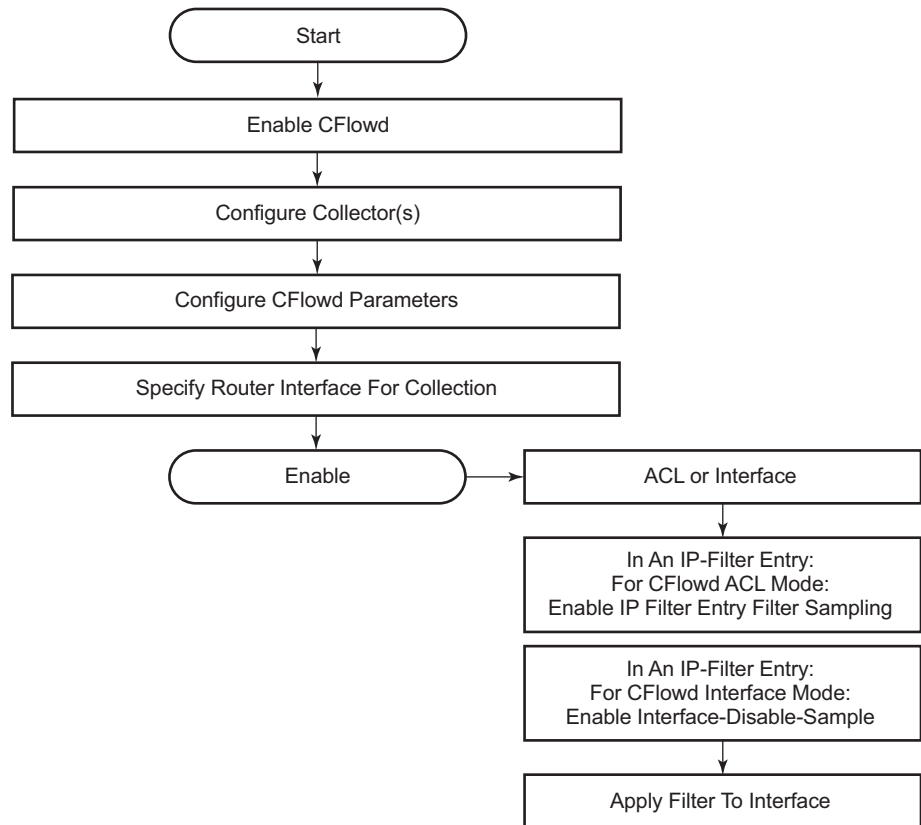
## 6.1.2 Cflowd Filter Matching

In the filter-matching process, usually every packet is matched against filter (access list) criteria to determine acceptability. With cflowd, only the first packet of a flow is checked. If the first packet is forwarded, an entry is added to the cflowd cache. Subsequent packets in the same flow are then forwarded without needing to be matched against the complete set of filters. Specific performance varies depending on the number and complexity of the filters.

## 6.2 Cflowd Configuration Process Overview

Figure 42 shows the process to configure cflowd parameters.

**Figure 42 Cflowd Configuration and Implementation Flow**



Router\_Config\_32

There are three modes in which cflowd can be enabled to sample traffic on an interface:

- Cflowd interface, where all traffic entering a specified port will be subjected to sampling at the configured sampling rate.
- Cflowd interface plus the definition of IP filters that specify an action of interface-disable-sample, where traffic that matches these filter entries will not be subject to cflowd sampling.
- Cflowd ACL, where IP filters must be created with entries containing the action filter-sampled. In this mode, only traffic matching these filter entries will be subject to the cflowd sampling process.



---

## 6.3 Configuration Notes

The following cflowd components must be configured for cflowd to be operational:

- Cflowd is enabled globally.
- At least one collector must be configured and enabled.
- A cflowd option must be specified and enabled on a router interface.
- Sampling must be enabled on either:
  - An IP filter that is applied to a port or service.
  - An interface on a port or service.

For the 7450 ESS, cflowd is only available when mixed-mode is enabled on the system.



---

## 6.4 Configuring Cflowd with CLI

This section provides information to configure cflowd using the command line interface.

### 6.4.1 Cflowd Configuration Overview

SR OS implementation of cflowd supports the option to analyze traffic flow. The implementation also supports the use of traffic/access list (ACL) filters to limit the type of traffic that is analyzed.

#### 6.4.1.1 Traffic Sampling

Traffic sampling does not examine all packets received by a router. Command parameters allow the rate at which traffic is sampled and sent for flow analysis to be modified. The default sampling rate is every 1000th packet. Excessive sampling over an extended period of time, for example, more than every 1000th packet, can burden router processing resources.

The following data is maintained for each individual flow in the raw flow cache:

- Source IP address
- Destinations IP address
- Source port
- Destination port
- Forwarding status
- Input interface
- Output interface
- IP protocol
- TCP flags
- First timestamp (of the first packet in the flow)
- Last timestamp (timestamp of last packet in the flow prior to expiry of the flow)
- Source AS number for peer and origin (taken from BGP)
- Destination AS number for peer and origin (taken from BGP)
- IP next hop
- BGP next hop

- ICMP type and code
- IP version
- Source prefix (from routing)
- Destination prefix (from routing)
- MPLS label stack from label 1 to 6

Within the raw flow cache, the following characteristics are used to identify an individual flow:

- Ingress interface
- Source IP address
- Destination IP address
- Source transport port number
- Destination transport port number
- IP protocol type
- IP TOS byte
- Virtual router ID
- ICMP type and code
- Direction
- MPLS labels

SR OS implementation allows you to enable cflowd either at the interface level or as an action to a filter. By enabling cflowd at the interface level, all IP packets forwarded by the interface are subject to cflowd analysis. By setting cflowd as an action in a filter, only packets matching the specified filter are subject to cflowd analysis. This provides the network operator greater flexibility in the types of flows that are captured.

### 6.4.1.2 Collectors

A collector defines how data flows should be exported from the flow cache. A maximum of five collectors can be configured. Each collector is identified by a unique IP address and UDP port value. Each collector can only export traffic in one version type: V5, V8, V9, or V10.

The parameters within a collector configuration can be modified or the defaults retained.

The **autonomous-system-type** command defines whether the autonomous system information to be included in the flow data is based on the originating AS or external peer AS of the flow.

#### 6.4.1.2.1 Aggregation

V8 aggregation allows for flow data to be aggregated into larger, less granular flows. Use aggregation commands to specify the type of data to be collected. These aggregation types are only applicable to flows being exported to a V8 collector.

The following aggregation schemes are supported:

- AS matrix — Flows are aggregated based on source and destination AS and ingress and egress interface.
- Protocol-port — Flows are aggregated based on the IP protocol, source port number, and destination port number.
- Source prefix — Flows are aggregated based on source prefix and mask, source AS, and ingress interface.
- Destination prefix — Flows are aggregated based on destination prefix and mask, destination AS, and egress interface.
- Source-destination prefix — Flows are aggregated based on source prefix and mask, destination prefix and mask, source and destination AS, ingress interface and egress interface.
- Raw — Flows are not aggregated and are sent to the collector in a V5 record.

### 6.4.2 Basic Cflowd Configuration

This section provides information to configure cflowd and examples of common configuration tasks. To sample traffic, the following parameters must be configured, as a minimum.

- Cflowd must be enabled.
- At least one collector must be configured and enabled.
- Sampling must be enabled on either:
  - An IP filter entry (and applied to a service or a port).
  - An interface applied to a port.

The following example shows a cflowd configuration:

```
A:ALA-1>config>cflowd# info detail
```

```

 active-timeout 30
 cache-size 65536
 inactive-timeout 15
 overflow 1
 rate 1000
 collector 10.10.10.103:2055 version 9
 no aggregation
 autonomous-system-type origin
 description "V9 collector"
 no shutdown
 exit
 template-retransmit 330
 exit
 no shutdown

A:ALA-1>config>cflowd#
```

## 6.4.3 Common Configuration Tasks

This section provides a brief overview of the tasks that must be performed to configure cflowd and provides the CLI commands. To begin traffic flow sampling, cflowd must be enabled and at least one collector must be configured.

### 6.4.3.1 Global Cflowd Components

The following common (global) attributes apply to all instances of cflowd:

- Active timeout — Controls the maximum time a flow record can be active before it will be automatically exported to defined collectors.
- Inactive timeout — Controls the minimum time before a flow is declared inactive. If no traffic is sampled for a flow for the inactive timeout duration, the flow is declared inactive and marked to be exported to the defined collectors.
- Cache size — Defines the maximum size of the flow cache.
- Overflow — Defines the percentage of flow records that are exported to all collectors if the flow cache size is exceeded.
- Rate — Defines the system-wide sampling rate for cflowd.
- Template retransmit— Defines the interval (in seconds) at which the V9 and V10 templates are retransmitted to all configured V9 or V10 collectors.

### 6.4.3.2 Enabling Cflowd

Cflowd is disabled by default. Cflowd is not shut down but must be configured, including at least one collector, to be active. Executing the **cflowd** command enables cflowd.

Use the following CLI syntax to enable cflowd:

**CLI Syntax:**    config# cflowd  
                   no shutdown

The following example shows the default values when cflowd is initially enabled. No collectors or collector options are configured.

```
A:ALA-1>config# info detail
...
#-----
echo "Cflowd Configuration"
#-----
 cflowd
 active-timeout 30
 cache-size 65536
 inactive-timeout 15
 overflow 1
 rate 1000
 template-retransmit 600
 no shutdown
 exit
#-----
A:ALA-1>config#
```

### 6.4.3.3 Configuring Global Cflowd Parameters

The following cflowd parameters apply to all instances where cflowd is enabled.

Use the following CLI commands to configure cflowd parameters:

**CLI Syntax:**    config>cflowd#  
                   active-timeout *minutes*  
                   cache-size *num-entries*  
                   enhanced-distribution  
                   export-mode {automatic | manual}  
                   inactive-timeout *seconds*  
                   overflow *percent*  
                   rate *sample-rate*  
                   template-retransmit *seconds*  
                   no shutdown

The following example shows a sample cflowd configuration:

```
A:ALA-1>config>cflowd# info
#-----
 active-timeout 20
 inactive-timeout 10
 overflow 10
 rate 100
#-----
A:ALA-1>config>cflowd#
```

### 6.4.3.4 Configuring Cflowd Collectors

To configure cflowd collector parameters, enter the following commands:

**CLI Syntax:**

```
config>cflowd#
collector ip-address[:port] [version version]
 aggregation
 as-matrix
 destination-prefix
 protocol-port
 raw
 source-destination-prefix
 source-prefix
 autonomous-system-type [origin | peer]
 description description-string
 no shutdown
 template-set {basic | mpls-ip | l2-ip | mpls-
 transport}
```

The following example shows a basic cflowd configuration:

```
A:ALA-1>config>cflowd# info

active-timeout 20
 inactive-timeout 10
 overflow 10
 rate 100
 collector 10.10.10.1:2000 version 8
 aggregation
 as-matrix
 raw
 exit
 description "AS info collector"
 exit
 collector 10.10.10.2:5000 version 8
 aggregation
 protocol-port
 source-destination-prefix
 exit
 autonomous-system-type peer
```



```

 description "Neighbor collector"
 exit

A:ALA-1>config>cflowd#

```

#### Version 9 collector example:

```

collector 10.10.10.9:2000 version 9
 description "v9collector"
 template-set mpls-ip
 no shutdown
exit

```

### 6.4.3.4.1 Version 9 and Version 10 Templates

If the collector is configured to use either V9 or V10 (IPFIX) formats, the flow data is sent to the designated collector using one of the predefined templates. The template used is based on the type of flow for which the data was collected (IPv4, IPv6, MPLS, or Ethernet (Layer 2)), and the configuration of the **template-set** parameter.

[Table 77](#) lists these values and the corresponding template used to export the flow data.

**Table 77**      **Template Sets**

| Traffic Flow | Basic      | MPLS-IP   |
|--------------|------------|-----------|
| IPv4         | Basic IPv4 | MPLS-IPv4 |
| IPv6         | Basic IPv6 | MPLS-IPv6 |
| MPLS         | Basic MPLS | MPLS-IP   |
| Ethernet     | L2-IP      | L2-IP     |

Each flow exported to a collector configured for either V9 or V10 formats will be sent using one of the flow template sets listed in [Table 77](#).

[Table 78](#) to [Table 85](#) list the fields in each template listed in [Table 77](#).

**Table 78**      **Basic IPv4 Template**

| Field Name    | Field ID |
|---------------|----------|
| IPv4 Src Addr | 8        |

**Table 78 Basic IPv4 Template (Continued)**

| Field Name                           | Field ID |
|--------------------------------------|----------|
| IPv4 Dest Addr                       | 12       |
| IPv4 Nexthop                         | 15       |
| BGP Nexthop                          | 18       |
| Ingress Interface                    | 10       |
| Egress Interface                     | 14       |
| Packet Count                         | 2        |
| Byte Count                           | 1        |
| Start Time                           | 22       |
| End Time                             | 21       |
| Flow Start Milliseconds <sup>1</sup> | 152      |
| Flow End Milliseconds <sup>1</sup>   | 153      |
| Src Port                             | 7        |
| Dest Port                            | 11       |
| Forwarding Status                    | 89       |
| TCP control Bits (Flags)             | 6        |
| IPv4 Protocol                        | 4        |
| IPv4 TOS                             | 5        |
| IP version                           | 60       |
| ICMP Type & Code                     | 32       |
| Direction                            | 61       |
| BGP Source ASN                       | 16       |
| BGP Dest ASN                         | 17       |
| Source IPv4 Prefix Length            | 9        |
| Dest IPv4 Prefix Length              | 13       |
| Minimum IP Total Length              | 25       |
| Maximum IP Total Length              | 26       |
| Minimum TTL                          | 52       |

**Table 78 Basic IPv4 Template (Continued)**

| Field Name                   | Field ID |
|------------------------------|----------|
| Maximum TTL                  | 53       |
| Multicast Replication Factor | 99       |
| IsMulticast <sup>1</sup>     | 206      |
| Ingress VRFID <sup>1</sup>   | 234      |
| Egress VRFID <sup>1</sup>    | 235      |

**Note:**

1. Only sent to collectors configured for V10 format.

**Table 79 MPLS-IPv4 Template**

| Field Name                           | Field ID |
|--------------------------------------|----------|
| IPv4 Src Addr                        | 8        |
| IPv4 Dest Addr                       | 12       |
| IPv4 Nexthop                         | 15       |
| BGP Nexthop                          | 18       |
| Ingress Interface                    | 10       |
| Egress Interface                     | 14       |
| Packet Count                         | 2        |
| Byte Count                           | 1        |
| Start Time                           | 22       |
| End Time                             | 21       |
| Flow Start Milliseconds <sup>1</sup> | 152      |
| Flow End Milliseconds <sup>1</sup>   | 153      |
| Src Port                             | 7        |
| Dest Port                            | 11       |
| Forwarding Status                    | 89       |
| TCP control Bits (Flags)             | 6        |
| IPv4 Protocol                        | 4        |

**Table 79**      **MPLS-IPv4 Template (Continued)**

| Field Name                   | Field ID |
|------------------------------|----------|
| IPv4 TOS                     | 5        |
| IP version                   | 60       |
| ICMP Type & Code             | 32       |
| Direction                    | 61       |
| BGP Source ASN               | 16       |
| BGP Dest ASN                 | 17       |
| Source IPv4 Prefix Length    | 9        |
| Dest IPv4 Prefix Length      | 13       |
| MPLS Top Label Type          | 46       |
| MPLS Top Label IPv4 Addr     | 47       |
| MPLS Label 1                 | 70       |
| MPLS Label 2                 | 71       |
| MPLS Label 3                 | 72       |
| MPLS Label 4                 | 73       |
| MPLS Label 5                 | 74       |
| MPLS Label 6                 | 75       |
| Minimum IP Total Length      | 25       |
| Maximum IP Total Length      | 26       |
| Minimum TTL                  | 52       |
| Maximum TTL                  | 53       |
| Multicast Replication Factor | 99       |
| IsMulticast <sup>1</sup>     | 206      |
| Ingress VRFID <sup>1</sup>   | 234      |
| Egress VRFID <sup>1</sup>    | 235      |

**Note:**

1. Only sent to collectors configured for V10 format.

**Table 80 Basic IPv6 Template**

| Field Name                           | Field ID |
|--------------------------------------|----------|
| IPv6 Src Addr                        | 27       |
| IPv6 Dest Addr                       | 28       |
| IPv6 Nexthop                         | 62       |
| IPv6 BGP Nexthop                     | 63       |
| IPv4 Nexthop                         | 15       |
| IPv4 BGP Nexthop                     | 18       |
| Ingress Interface                    | 10       |
| Egress Interface                     | 14       |
| Packet Count                         | 2        |
| Byte Count                           | 1        |
| Start Time                           | 22       |
| End Time                             | 21       |
| Flow Start Milliseconds <sup>1</sup> | 152      |
| Flow End Milliseconds <sup>1</sup>   | 153      |
| Src Port                             | 7        |
| Dest Port                            | 11       |
| Forwarding Status                    | 89       |
| TCP control Bits (Flags)             | 6        |
| Protocol                             | 4        |
| IPv6 Extension Hdr                   | 64       |
| IPv6 Next Header                     | 193      |
| IPv6 Flow Label                      | 31       |
| TOS                                  | 5        |
| IP version                           | 60       |
| IPv6 ICMP Type & Code                | 139      |
| Direction                            | 61       |
| BGP Source ASN                       | 16       |

**Table 80 Basic IPv6 Template (Continued)**

| Field Name                   | Field ID |
|------------------------------|----------|
| BGP Dest ASN                 | 17       |
| IPv6 Src Mask                | 29       |
| IPv6 Dest Mask               | 30       |
| Minimum IP Total Length      | 25       |
| Maximum IP Total Length      | 26       |
| Minimum TTL                  | 52       |
| Maximum TTL                  | 53       |
| Multicast Replication Factor | 99       |
| IsMulticast <sup>1</sup>     | 206      |
| Ingress VRFID <sup>1</sup>   | 234      |
| Egress VRFID <sup>1</sup>    | 235      |

**Note:**

1. Only sent to collectors configured for V10 format.

**Table 81 MPLS-IPv6 Template**

| Field Name        | Field ID |
|-------------------|----------|
| IPv6 Src Addr     | 27       |
| IPv6 Dest Addr    | 28       |
| IPv6 Nexthop      | 62       |
| IPv6 BGP Nexthop  | 63       |
| IPv4 Nexthop      | 15       |
| IPv4 BGP Nexthop  | 18       |
| Ingress Interface | 10       |
| Egress Interface  | 14       |
| Packet Count      | 2        |
| Byte Count        | 1        |
| Start Time        | 22       |

**Table 81 MPLS-IPv6 Template (Continued)**

| Field Name                           | Field ID |
|--------------------------------------|----------|
| End Time                             | 21       |
| Flow Start Milliseconds <sup>1</sup> | 152      |
| Flow End Milliseconds <sup>1</sup>   | 153      |
| Src Port                             | 7        |
| Dest Port                            | 11       |
| Forwarding Status                    | 89       |
| TCP control Bits (Flags)             | 6        |
| Protocol                             | 4        |
| IPv6 Extension Hdr                   | 64       |
| IPv6 Next Header                     | 193      |
| IPv6 Flow Label                      | 31       |
| TOS                                  | 5        |
| IP version                           | 60       |
| IPv6 ICMP Type & Code                | 139      |
| Direction                            | 61       |
| BGP Source ASN                       | 16       |
| BGP Dest ASN                         | 17       |
| IPv6 Src Mask                        | 29       |
| IPv6 Dest Mask                       | 30       |
| MPLS Top Label Type                  | 46       |
| MPLS Top Label IPv6 Addr             | 47       |
| MPLS Label 1                         | 70       |
| MPLS Label 2                         | 71       |
| MPLS Label 3                         | 72       |
| MPLS Label 4                         | 73       |
| MPLS Label 5                         | 74       |
| MPLS Label 6                         | 75       |

**Table 81 MPLS-IPv6 Template (Continued)**

| Field Name                   | Field ID |
|------------------------------|----------|
| MPLS_TOP_LABEL_TYPE          | 46       |
| MPLS_TOP_LABEL_ADDR          | 47       |
| Minimum IP Total Length      | 25       |
| Maximum IP Total Length      | 26       |
| Minimum TTL                  | 52       |
| Maximum TTL                  | 53       |
| Multicast Replication Factor | 99       |
| IsMulticast <sup>1</sup>     | 206      |
| Ingress VRFID <sup>1</sup>   | 234      |
| Egress VRFID <sup>1</sup>    | 235      |

**Note:**

1. Only sent to collectors configured for V10 format.

**Table 82 Basic MPLS Template**

| Field Name                           | Field ID |
|--------------------------------------|----------|
| Start Time                           | 22       |
| End Time                             | 21       |
| Flow Start Milliseconds <sup>1</sup> | 152      |
| Flow End Milliseconds <sup>1</sup>   | 153      |
| Ingress Interface                    | 10       |
| Egress Interface                     | 14       |
| Packet Count                         | 2        |
| Byte Count                           | 1        |
| Direction                            | 61       |
| MPLS Top Label Type                  | 46       |
| MPLS Top Label Address               | 47       |
| MPLS Label 1                         | 70       |



**Table 82 Basic MPLS Template (Continued)**

| Field Name   | Field ID |
|--------------|----------|
| MPLS Label 2 | 71       |
| MPLS Label 3 | 72       |
| MPLS Label 4 | 73       |
| MPLS Label 5 | 74       |
| MPLS Label 6 | 75       |

**Note:**

1. Only sent to collectors configured for V10 format.

**Table 83 MPLS-IP Template**

| Field Name                           | Field ID |
|--------------------------------------|----------|
| IPv4 Src Addr                        | 8        |
| IPv4 Dest Addr                       | 12       |
| IPv4 Nexthop                         | 15       |
| IPv6 Src Addr                        | 27       |
| IPv6 Dest Addr                       | 28       |
| IPv6 Nexthop                         | 62       |
| Ingress Interface                    | 10       |
| Egress Interface                     | 14       |
| Packet Count                         | 2        |
| Byte Count                           | 1        |
| Start Time                           | 22       |
| End Time                             | 21       |
| Flow Start Milliseconds <sup>1</sup> | 152      |
| Flow End Milliseconds <sup>1</sup>   | 153      |
| Src Port                             | 7        |
| Dest Port                            | 11       |
| TCP control Bits (Flags)             | 6        |

**Table 83 MPLS-IP Template (Continued)**

| Field Name               | Field ID |
|--------------------------|----------|
| IPv4 Protocol            | 4        |
| IPv4 TOS                 | 5        |
| IP version               | 60       |
| ICMP Type & Code         | 32       |
| Direction                | 61       |
| MPLS Top Label Type      | 46       |
| MPLS Top Label IPv4 Addr | 47       |
| MPLS Label 1             | 70       |
| MPLS Label 2             | 71       |
| MPLS Label 3             | 72       |
| MPLS Label 4             | 73       |
| MPLS Label 5             | 74       |
| MPLS Label 6             | 75       |

**Note:**

1. Only sent to collectors configured for V10 format.

**Table 84 L2-IP (Ethernet) Flow Template**

| Field Name <sup>1</sup>     | Field ID |
|-----------------------------|----------|
| MAC Src Addr                | 56       |
| MAC Dest Addr               | 80       |
| Ingress Physical Interface  | 252      |
| Egress Physical Interface   | 253      |
| Dot1q VLAN ID               | 243      |
| Dot1q Customer VLAN ID      | 245      |
| Post Dot1q VLAN ID          | 254      |
| Post Dot1q Customer VLAN Id | 255      |
| IPv4 Src Addr               | 8        |

**Table 84 L2-IP (Ethernet) Flow Template (Continued)**

| Field Name <sup>1</sup>  | Field ID |
|--------------------------|----------|
| IPv4 Dest Addr           | 12       |
| IPv6 Src Addr            | 27       |
| IPv6 Dest Addr           | 28       |
| Packet Count             | 2        |
| Byte Count               | 1        |
| Flow Start Milliseconds  | 152      |
| Flow End Milliseconds    | 153      |
| Src Port                 | 7        |
| Dest Port                | 11       |
| TCP control Bits (Flags) | 6        |
| Protocol                 | 4        |
| IPv6 Option Header       | 64       |
| IPv6 Next Header         | 196      |
| IPv6 Flow Label          | 31       |
| TOS                      | 5        |
| IP Version               | 60       |
| ICMP Type Code           | 32       |

**Note:**

1. Only one Ethernet (L2-IP) flow template is supported and exported to IPFIX (V10) collectors.

**Table 85 MPLS-IP Template**

| Field Name              | Field ID |
|-------------------------|----------|
| Flow Start Milliseconds | 152      |
| Flow End Milliseconds   | 153      |
| VRF ID                  | 234      |
| Ingress Interface       | 10       |
| Packet Count            | 2        |

**Table 85 MPLS-IP Template (Continued)**

| Field Name          | Field ID |
|---------------------|----------|
| Byte Count          | 1        |
| Direction           | 61       |
| MPLS_TOP_LABEL_TYPE | 46       |
| MPLS_TOP_LABEL_ADDR | 47       |
| MPLS Label-1        | 70       |

### 6.4.3.5 Specifying Cflowd Options on an IP Interface

When cflowd is enabled on an interface, all packets forwarded by the interface are subject to analysis according to the global cflowd configuration and sorted according to the collector configurations.

See [Table 86](#) for configuration combinations.

When the **cflowd-parameters sampling unicast type interface** option is configured in the **config>router>interface** context, the following requirements must be met to enable traffic sampling on the interface:

- Cflowd must be enabled.
- At least one cflowd collector must be configured and enabled.
- The **config>router>interface>cflowd-parameters sampling {unicast | multicast} type {acl | interface} [direction {ingress-only | egress-only | both}]** must be performed. By default, the direction is ingress-only.
- To prevent certain types of traffic from being sampled when interface sampling is enabled, use the **interface-disable-sample** command in the **config>filter>ip-filter** or **config>filter>ipv6-filter** context. The filter must be applied to the service or network interface on which the traffic to be omitted is to ingress the system.

#### 6.4.3.5.1 Interface Configurations

**CLI Syntax:**

```
config>router>if#
cflowd-parameters
 sampling {unicast | multicast} type {acl |
interface}
```

```
no sampling {unicast | multicast} type
{acl|interface}
```

Depending on the option selected, either **acl** or **interface**, cflowd extracts traffic flow samples from an IP filter or an interface for analysis. All packets forwarded by the interface are analyzed according to the cflowd configuration.

The **acl** option must be selected to enable traffic sampling on an IP filter. Cflowd (**filter-sample**) must be enabled in at least one IP filter entry.

The **interface** option must be selected to enable traffic sampling on an interface. If cflowd is not enabled (**no cflowd**), traffic sampling will not occur on the interface.

#### 6.4.3.5.2 Service Interfaces

**CLI Syntax:** `config>router>interface# cflowd-parameters sampling {unicast | multicast} type {acl | interface}`

When enabled on a service interface, cflowd collects routed traffic flow samples through a router for analysis. Cflowd is supported on IES and VPRN services interfaces only. Layer 2 traffic is excluded. All packets forwarded by the interface are analyzed according to the cflowd configuration. On the interface level, cflowd can be associated with a filter (ACL) or an IP interface. L2 cflowd ingress sampling is supported on VPLS and Epipe SAPs.

#### 6.4.3.6 Specifying Sampling Options in Filter Entries

Packets are matched against filter entries to determine acceptability. With cflowd, only the first packet of a flow is compared. If the first packet matches the filter criteria, an entry is added to the cflowd cache. Subsequent packets in the same flow are also sampled based on the cache entry.

Because a filter can be applied to more than one interface (when configured with a **scope template**), the **interface-disable-sample** option is intended to enable or disable traffic sampling on an interface-by-interface basis. The command can be enabled or disabled as needed instead of having to create numerous filter versions.

To enable an interface for filter traffic sampling, the following requirements must be met:

- Cflowd must be enabled globally.
- At least one cflowd collector must be configured and enabled.

- On the IP interface being used, the **config>router>interface>cflowd-parameters sampling {unicast | multicast} type acl** option must be selected.
- On the IP filter being used, the **entry>filter-sample** option must be explicitly enabled for the entries matching the traffic that should be sampled. The default is **no filter-sample**. See Filter Configuration for more information.
- The filter must be applied to a service or a network interface. The service or port must be enabled and operational.

#### 6.4.3.6.1 Filter Configurations

**CLI Syntax:**     config>filter>ip-filter>entry#  
                  [no] filter-sample  
                  [no] interface-disable-sample

or

**CLI Syntax:**     config>filter>ipv6-filter>entry#  
                  [no] filter-sample  
                  [no] interface-disable-sample

When a filter policy is applied to a service or a network interface, sampling can be configured so that traffic matching the associated IP filter entry is sampled when the IP interface is set to cflowd ACL mode and the **filter-sample** command is enabled. If cflowd is either not enabled (**no filter-sample**) or set to the cflowd interface mode, sampling does not occur.

When the **interface-disable-sample** command is enabled, traffic matching the associated IP filter entry is not sampled if the IP interface is set to cflowd ACL mode.

#### 6.4.3.6.2 Dependencies

For cflowd to be operational, the following requirements must be met:

- Cflowd must be enabled on a global level. If cflowd is disabled, any traffic sampling instances are also disabled.
- At least one collector must be configured and enabled in order for traffic sampling to occur on an enabled entity.
- If a specific collector UDP port is not identified, flows are sent to port 2055 by default.

Cflowd can also be dependent on the following entity configurations:

- [Interface Configurations](#)
- [Service Interfaces](#)
- [Filter Configurations](#)

The combination of interface and filter entry configurations determines whether flow sampling occurs. [Table 86](#) lists the expected results based on cflowd configuration dependencies.

**Table 86 Cflowd Configuration Dependencies**

| Interface Setting                                 | cflowd-parameter type Setting | Command ip-filter entry Setting | Expected Results                                                                             |
|---------------------------------------------------|-------------------------------|---------------------------------|----------------------------------------------------------------------------------------------|
| IP-filter mode                                    | ACL                           | filter-sample                   | Traffic matching is sampled at specified rate.                                               |
| IP-filter mode                                    | ACL                           | no filter-sample                | No traffic is sampled on this interface.                                                     |
| IP-filter mode or cflowd not enabled on interface | ACL                           | interface-disable-sample        | Command is ignored. No sampling occurs.                                                      |
| Interface mode                                    | Interface                     | interface-disable-sample        | Traffic matching this IP filter entry is not sampled.                                        |
| Interface mode                                    | Interface                     | none                            | All IP traffic ingressing the interface is subject to sampling.                              |
| Interface mode                                    | Interface                     | filter-sample                   | Filter-level action is ignored. All traffic ingressing the interface is subject to sampling. |

## 6.5 Cflowd Configuration Management Tasks

This section describes Cflowd configuration management tasks:

### 6.5.1 Modifying Global Cflowd Components

Cflowd parameter modifications apply to all instances where cflowd is enabled. Changes are applied immediately. Use the following cflowd commands to modify global cflowd parameters:

**CLI Syntax:**

```
config>cflowd#
active-timeout minutes
no active-timeout
cache-size num-entries
no cache-size
[no] enhanced-distribution
export-mode {automatic | manual}
inactive-timeout seconds
no inactive-timeout
overflow percent
no overflow
rate sample-rate
no rate
[no] shutdown
template-retransmit seconds
no template-retransmit
```

The following example shows the cflowd command syntax to modify configuration parameters:

**Example:**

```
config>cflowd# active-timeout 60
config>cflowd# no inactive-timeout
config>cflowd# overflow 2
config>cflowd# rate 10
```

The following example shows an example cflowd configuration:

```
A:ALA-1>config>cflowd# info
#-----
 active-timeout 60
 overflow 2
 rate 10
#-----
A:ALA-1>config>cflowd#
```



## 6.5.2 Modifying Cflowd Collector Parameters

Use the following commands to modify cflowd collector and aggregation parameters:

**CLI Syntax:**

```
config>cflowd#
collector ip-address[:port] [version version]
no collector ip-address[:port]
 [no] aggregation
 [no] as-matrix
 [no] destination-prefix
 [no] protocol-port
 [no] raw
 [no] source-destination-prefix
 [no] source-prefix
 [no] autonomous-system-type [origin | peer]
 [no] description description-string
 [no] shutdown
 template-set {basic | mpls-ip | l2-ip | mpls-
 transport}
```

If a specific collector UDP port is not identified, flows are sent to port 2055 by default.

The following example displays basic cflowd modifications:

```
A:ALA-1>config>cflowd# info

 active-timeout 60
 overflow 2
 rate 10
 collector 10.10.10.1:2000 version 5
 description "AS info collector"
 exit
 collector 10.10.10.2:5000 version 8
 aggregation
 source-prefix
 raw
 exit
 description "Test collector"
 exit

A:ALA-1>config>cflowd#
```



## 6.6 Cflowd Configuration Command Reference

- [Command Hierarchies](#)
- [Command Descriptions](#)

### 6.6.1 Command Hierarchies

```
config
— [no] cflowd
 — active-timeout minutes
 — no active-timeout
 — cache-size num-entries
 — no cache-size
 — collector ip-address[:port] [version {5 | 8 | 9 | 10}]
 — no collector ip-address[:port]
 — [no] aggregation
 — [no] as-matrix
 — [no] destination-prefix
 — [no] protocol-port
 — [no] raw
 — [no] source-destination-prefix
 — [no] source-prefix
 — autonomous-system-type {origin | peer}
 — no autonomous-system-type
 — description description-string
 — no description
 — export-filter
 — [no] family
 — [no] ipv4
 — [no] ipv6
 — [no] l2-ip
 — [no] mcast-ipv4
 — [no] mcast-ipv6
 — [no] mpls
 — [no] router {router-name | vprn-svc-id}
 — router {router-name | vprn-svc-id}
 — [no] shutdown
 — template-set {basic | mpls-ip | l2-ip | mpls-transport}
 — [no] enhanced-distribution
 — export-mode [automatic | manual]
 — inactive-timeout seconds
 — no inactive-timeout
 — overflow percent
 — no overflow
 — rate sample-rate
 — no rate
 — [no] shutdown
 — template-retransmit seconds
```

- no **template-retransmit**
- [no] **use-vrtr-if-index**

## 6.6.2 Command Descriptions

### 6.6.2.1 Global Commands

#### cflowd

|                    |                                                                                                                                                                                                                                                                |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [no] <b>cflowd</b>                                                                                                                                                                                                                                             |
| <b>Context</b>     | config>cflowd                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command creates the context to configure cflowd.</p> <p>The <b>no</b> form of this command removes all configuration under cflowd including the deletion of all configured collectors. This can only be executed if cflowd is in a shutdown state.</p> |
| <b>Default</b>     | no cflowd                                                                                                                                                                                                                                                      |

#### active-timeout

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>active-timeout</b> <i>minutes</i><br><b>no active-timeout</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>cflowd                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Description</b> | <p>This command configures the maximum amount of time before an active flow is aged out of the active cache. If an individual flow is active for this amount of time, the flow is aged out and a new flow will be created on the next packet sampled for that flow.</p> <p>Existing flows do not inherit the new active-timeout value if this parameter is changed while cflowd is active. The active-timeout value for a flow is set when the flow is first created in the active cache table and does not change dynamically.</p> <p>The <b>no</b> form of this command resets the inactive timeout back to the default value.</p> |
| <b>Default</b>     | active-timeout 30                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Parameters</b>  | <i>minutes</i> — the value expressed in minutes before an active flow is exported                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Values</b>      | 1 to 600                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## cache-size

|                    |                                                                                                                                                                                                         |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>cache-size</b> <i>num-entries</i><br><b>no cache-size</b>                                                                                                                                            |
| <b>Context</b>     | config>cflowd                                                                                                                                                                                           |
| <b>Description</b> | This command specifies the maximum number of active flows to maintain in the flow cache table.<br><br>The <b>no</b> form of this command resets the number of active entries back to the default value. |
| <b>Default</b>     | cache-size 65536                                                                                                                                                                                        |
| <b>Parameters</b>  | <i>num-entries</i> — Specifies the maximum number of entries maintained in the cflowd cache. The number depends on the CPM version.                                                                     |

### Values

For the 7450 ESS and 7750 SR 1000 to 250000  
(cfm-xp, SF/CPM3):

For the 7450 ESS and 7750 SR 1000 to 1000000  
(CPM4 or CPM5):

For the 7950 XRS: 1000 to 1500000

### Default

For the 7450 ESS and 7750 SR: 65536 (64K)

For the 7950 XRS: 500000

## collector

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>collector</b> <i>ip-address[:port]</i> { <b>version</b> [5   8   9   10]}<br><b>no collector</b> <i>ip-address[:port]</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Context</b>     | config>cflowd                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b> | This command defines a flow data collector for cflowd data. The IP address of the flow collector must be specified. The UDP port number is an optional parameter. If it is not set, the default of 2055 is used for all collector versions. To connect to a IPFIX (version 10) collector using the IPFIX default port, specify port 4739 when configuring the collector. The version must be specified. A maximum of 5 collectors can be configured.<br><br>The <b>no</b> form of this command removes the flow collector definition from the config and stops the export of data to the collector. The collector needs to be shut down to be deleted. |

|                   |                                                                                                                                                                            |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Parameters</b> | <i>ip-address</i> — Specifies the address of a remote cflowd collector host to receive the exported cflowd data.                                                           |
|                   | <b>Values</b>                                                                                                                                                              |
|                   | <div>&lt;ip-address[:port]&gt;: ip-address - a.b.c.d[:port] (IPv4)</div> <div>x::x:x:x:x:x:x (IPv6)</div> <div>[x:x:x:x:x:x:x]:port (IPv6)</div> <div>x - [0..FFFF]H</div> |
|                   | <i>port</i> — Specifies the UDP port number on the remote cflowd collector host to receive the exported cflowd data.                                                       |
|                   | <b>Values</b> 1 to 65535                                                                                                                                                   |
|                   | <b>Default</b> 2055                                                                                                                                                        |
|                   | <i>version</i> — Specifies the version of the flow data collector.                                                                                                         |
|                   | <b>Values</b> 5, 8, 9, 10                                                                                                                                                  |
|                   | <b>Default</b> 5                                                                                                                                                           |

aggregation

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] aggregation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>cflowd>collector                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b> | <p>This command configures the type of aggregation scheme to be exported.</p> <p>Specifies the type of data to be aggregated and to the collector.</p> <p>To configure aggregation, you must decide which type of aggregation scheme to configure: autonomous system, destination prefix, protocol port, raw, source destination, or source prefix.</p> <p>This can only be configured if the collector version is configured as V8.</p> <p>The <b>no</b> form of this command removes all aggregation types from the collector configuration.</p> |
| <b>Default</b>     | no aggregation                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

as-matrix

|                |                                     |
|----------------|-------------------------------------|
| <b>Syntax</b>  | <b>[no] as-matrix</b>               |
| <b>Context</b> | config>cflowd>collector>aggregation |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command specifies that the aggregation data should be based on autonomous system (AS) information. An AS matrix contains packet and byte counters for traffic from either source-destination autonomous systems or last-peer to next-peer autonomous systems.</p> <p>The <b>no</b> form of this command removes this type of aggregation from the collector configuration.</p> |
| <b>Default</b>     | no as-matrix                                                                                                                                                                                                                                                                                                                                                                           |

## destination-prefix

|                    |                                                                                                                                                                                                  |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] destination-prefix</b>                                                                                                                                                                   |
| <b>Context</b>     | config>cflowd>collector>aggregation                                                                                                                                                              |
| <b>Description</b> | <p>This command specifies that the aggregation data is based on destination prefix information.</p> <p>The <b>no</b> form removes this type of aggregation from the collector configuration.</p> |

## protocol-port

|                    |                                                                                                                                                                                                                                                |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] protocol-port</b>                                                                                                                                                                                                                      |
| <b>Context</b>     | config>cflowd>collector>aggregation                                                                                                                                                                                                            |
| <b>Description</b> | <p>This command specifies that flows be aggregated based on the IP protocol, source port number, and destination port number.</p> <p>The <b>no</b> form of this command removes this type of aggregation from the collector configuration.</p> |

## raw

|                    |                                                                                                                                                                                                   |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] raw</b>                                                                                                                                                                                   |
| <b>Context</b>     | config>cflowd>collector>aggregation                                                                                                                                                               |
| <b>Description</b> | <p>This command configures raw (unaggregated) flow data to be sent in Version 5.</p> <p>The <b>no</b> form of this command removes this type of aggregation from the collector configuration.</p> |
| <b>Default</b>     | n/a                                                                                                                                                                                               |

---

## source-destination-prefix

|                    |                                                                                                                                                                                                          |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] source-destination-prefix</b>                                                                                                                                                                    |
| <b>Context</b>     | config>cflowd>collector>aggregation                                                                                                                                                                      |
| <b>Description</b> | <p>This command configures cflowd aggregation based on source and destination prefixes.</p> <p>The <b>no</b> form of this command removes this type of aggregation from the collector configuration.</p> |
| <b>Default</b>     | n/a                                                                                                                                                                                                      |

## source-prefix

|                    |                                                                                                                                                                                                    |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] source-prefix</b>                                                                                                                                                                          |
| <b>Context</b>     | config>cflowd>collector>aggregation                                                                                                                                                                |
| <b>Description</b> | <p>This command configures cflowd aggregation based on source prefix information.</p> <p>The <b>no</b> form of this command removes this type of aggregation from the collector configuration.</p> |
| <b>Default</b>     | n/a                                                                                                                                                                                                |

## autonomous-system-type

|                    |                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>autonomous-system-type {origin   peer}</b><br><b>no autonomous-system-type</b>                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>cflowd>collector                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b> | <p>This command defines whether the autonomous system (AS) information included in the flow data is based on the originating AS or external peer AS of the routes.</p> <p>This option is only allowed if the collector is configured as Version 5 or Version 8.</p> <p>The <b>no</b> form of this command resets the AS type to the default value.</p> |
| <b>Default</b>     | autonomous-system-type origin                                                                                                                                                                                                                                                                                                                          |
| <b>Parameters</b>  | <p><b>origin</b> — Specifies that the AS information included in the flow data is based on the originating AS.</p> <p><b>peer</b> — Specifies that the AS information included in the flow data is based on the peer AS.</p>                                                                                                                           |



---

## description

|                    |                                                                                                                                                                                                                                                                                               |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>description</b> <i>description-string</i><br><b>no description</b>                                                                                                                                                                                                                         |
| <b>Context</b>     | config>cflowd>collector                                                                                                                                                                                                                                                                       |
| <b>Description</b> | <p>This command creates a text description stored in the configuration file for a configuration context.</p> <p>The <b>no</b> form of this command removes the description string from the context.</p>                                                                                       |
| <b>Parameters</b>  | <i>description-string</i> — the description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, \$, spaces, etc.), the entire string must be enclosed within double quotes. |

## export-filter

|                    |                                                                                                                                                                                                      |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>export-filter</b>                                                                                                                                                                                 |
| <b>Context</b>     | config>cflowd>collector                                                                                                                                                                              |
| <b>Description</b> | <p>This command creates the CLI context to specify cflowd data filters. These filters allow the administrator to control which flows are sent or are not sent to an associated cflowd collector.</p> |

## family

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                        |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] family</b>                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Context</b>     | config>cflowd>collector>export-filter                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command defines the address family for the flow types that should not be sent to the associated cflowd collector.</p> <p>Multiple family types can be defined in this context to filter out multiple address families to a given collector.</p> <p>The <b>no</b> form of this command removes the address family definition, allowing all address family types to be exported to the associated collector.</p> |
| <b>Default</b>     | no family                                                                                                                                                                                                                                                                                                                                                                                                              |

## ipv4

|               |                  |
|---------------|------------------|
| <b>Syntax</b> | <b>[no] ipv4</b> |
|---------------|------------------|

---

|                    |                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Context</b>     | config>cflowd>collector>export-filter>family                                                                                                                                                                          |
| <b>Description</b> | <p>This command filters IPv4 flow data from being sent to the associated collector.</p> <p>The <b>no</b> form of this command removes the filter, allowing IPv4 flow data to be sent to the associated collector.</p> |
| <b>Default</b>     | no ipv4                                                                                                                                                                                                               |

## ipv6

|                    |                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [no] ipv6                                                                                                                                                                                                             |
| <b>Context</b>     | config>cflowd>collector>export-filter>family                                                                                                                                                                          |
| <b>Description</b> | <p>This command filters IPv6 flow data from being sent to the associated collector.</p> <p>The <b>no</b> form of this command removes the filter, allowing IPv6 flow data to be sent to the associated collector.</p> |
| <b>Default</b>     | no ipv6                                                                                                                                                                                                               |

## l2-ip

|                    |                                                                                                                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [no] l2-ip                                                                                                                                                                                                                        |
| <b>Context</b>     | config>cflowd>collector>export-filter>family                                                                                                                                                                                      |
| <b>Description</b> | <p>This command filters Layer 2 IP flow data from being sent to the associated collector.</p> <p>The <b>no</b> form of this command removes the filter, allowing Layer 2 IP flow data to be sent to the associated collector.</p> |
| <b>Default</b>     | no l2-ip                                                                                                                                                                                                                          |

## mcast-ipv4

|                    |                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | [no] mcast-ipv4                                                                                                                                                                                                                           |
| <b>Context</b>     | config>cflowd>collector>export-filter>family                                                                                                                                                                                              |
| <b>Description</b> | <p>This command filters multicast IPv4 flow data from being sent to the associated collector.</p> <p>The <b>no</b> form of this command removes the filter, allowing multicast IPv4 flow data to be sent to the associated collector.</p> |
| <b>Default</b>     | no mcast-ipv4                                                                                                                                                                                                                             |

---

## mcast-ipv6

|                    |                                                                                                                                                                                                                                           |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] mcast-ipv6</b>                                                                                                                                                                                                                    |
| <b>Context</b>     | config>cflowd>collector>export-filter>family                                                                                                                                                                                              |
| <b>Description</b> | <p>This command filters multicast IPv6 flow data from being sent to the associated collector.</p> <p>The <b>no</b> form of this command removes the filter, allowing multicast IPv6 flow data to be sent to the associated collector.</p> |
| <b>Default</b>     | no mcast-ipv6                                                                                                                                                                                                                             |

## mpls

|                    |                                                                                                                                                                                                                       |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] mpls</b>                                                                                                                                                                                                      |
| <b>Context</b>     | config>cflowd>collector>export-filter>family                                                                                                                                                                          |
| <b>Description</b> | <p>This command filters MPLS flow data from being sent to the associated collector.</p> <p>The <b>no</b> form of this command removes the filter, allowing MPLS flow data to be sent to the associated collector.</p> |
| <b>Default</b>     | no mpls                                                                                                                                                                                                               |

## router

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] router</b> { <i>router-name</i>   <i>vprn-svc-id</i> }                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Context</b>     | config>cflowd>collector>export-filter                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b> | <p>This command allows the flow data from only specific router instances to be sent to the associated collector.</p> <p>Multiple router instances can be configured by issuing the command multiple times with the different router-instances.</p> <p>The <b>no</b> form of this command removes the specified router-instance restriction, which means flows from that router-instance will no longer be exported. If all router-instances are removed, then flows from all router instances are sent to the associated collector.</p> |
| <b>Default</b>     | no router                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Parameters</b>  | <i>router-name</i> — Specifies the router name. Only “Base” is supported.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Values</b>      | Base                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

---

*vprn-svc-id* — Specifies the router instance VPRN service ID.

**Values** 1 to 2147483647

## router

**Syntax** **router** {*router-name* | *vprn-svc-id*}

**Context** config>cflowd>collector

**Description** This command configures the flow data sent to the associated collector to be sent within the specified router context. If this parameter is not specified, flow data is exported using the management routing context.

**Default** router management

**Parameters** *router-name* — Specifies the router name.

**Values** Base, management

**Default** management

*vprn-svc-id* — Specifies the router instance VPRN service ID.

**Values** 1 to 2147483647

## shutdown

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

**Context** config>cflowd  
config>cflowd>collector

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

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

The **no** form of this command administratively enables an entity.

Unlike other commands and parameters where the default state is not indicated in the configuration file. The **shutdown** and **no shutdown** states are always indicated in system generated configuration files.

**Default** no shutdown

## template-set

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>template-set {basic   mpls-ip   l2-ip   mpls-transport}</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Context</b>     | config>cflowd>collector                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b> | This command specifies the set of templates sent to the collector when using cflowd Version 9 or Version 10.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Default</b>     | template-set basic                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Parameters</b>  | <p><b>basic</b> — Specifies that basic flow data is sent.</p> <p><b>mpls-ip</b> — Specifies that extended flow data is sent that includes IP and MPLS flow information.</p> <p><b>l2-ip</b> — Specifies that extended flow data is sent that includes Layer 2 (Ethernet) and IP flow information. This template is only applicable for V10 (IPFIX) collectors.</p> <p><b>mpls-transport</b> — Specifies that cflowd can collect flow statistics for MPLS traffic using only the outer transport label, EXP bit value, and ingress interface as the flow identifier. This template enables the collection of flow statistics on a core router to develop LSP usage statistics.</p> |

## enhanced-distribution

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>enhanced-distribution</b><br><b>no enhanced-distribution</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Context</b>     | config>cflowd                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b> | <p>This command enables the inclusion of the ingress port ID into the hash algorithm used to distribute cflowd sample traffic to cflowd processes running on the 7950 XRS CPM. By including this new attribute, cflowd may see better distribution of flows across processing tasks if there is a limited number of IP interfaces on which sampling is performed, but those interfaces use LAGs with a large number of port members.</p> <p>By enabling this option, the same flow may be captured multiple times if packets are received on multiple ingress ports.</p> <p>This command is only applicable to cflowd running on a 7950 XRS platform.</p> <p>The <b>no</b> form of this command removes the command from the configuration and disables the inclusion of the ingress port ID in the cflowd hash algorithm.</p> |
| <b>Default</b>     | no enhanced-distribution                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |

---

## export-mode

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>export-type</b> [automatic   manual]                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Context</b>     | config>cflowd                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | This command can be used to control how exports are generated by the cflowd process. The default behavior is for flow data to be exported automatically based on the active and inactive time-out values. The alternative mode is manual in which case flow data is only exported when the command “tools perform cflowd manual-export” is issued. The only exception is if the cflowd cache overflows, in which case the normal automatic export process is used. |
| <b>Default</b>     | export-mode automatic                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Parameters</b>  | <b>automatic</b> — cflowd flow data is automatically generated<br><b>manual</b> — cflowd flow data is exported only when manually triggered                                                                                                                                                                                                                                                                                                                        |

## inactive-timeout

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>inactive-timeout</b> <i>seconds</i><br><b>no inactive-timeout</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Context</b>     | config>cflowd                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command specifies the amount of time, in seconds, that must elapse without a packet matching a flow in order for the flow to be considered inactive.</p> <p>The <b>no</b> form of this command resets the inactive timeout back to the default of 15 seconds.</p> <p>Existing flows will not inherit the new inactive-timeout value if this parameter is changed while cflowd is active. The inactive-timeout value for a flow is set when the flow is first created in the active cache table and does not change dynamically.</p> |
| <b>Default</b>     | inactive-timeout 15                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Parameters</b>  | <b>seconds</b> — Specifies the amount of time, in seconds, that must elapse without a packet matching a flow in order for the flow to be considered inactive.<br><b>Values</b> 10 to 600                                                                                                                                                                                                                                                                                                                                                    |

## overflow

|                |                                                      |
|----------------|------------------------------------------------------|
| <b>Syntax</b>  | <b>overflow</b> <i>percent</i><br><b>no overflow</b> |
| <b>Context</b> | config>cflowd                                        |

---

|                    |                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b> | <p>This command specifies the percentage of the flow cache entries removed when the maximum number of entries is exceeded. The entries removed are the entries that have not been updated for the longest amount of time.</p> <p>The <b>no</b> form of this command resets the number of entries cleared from the flow cache on overflow to the default value.</p> |
| <b>Default</b>     | overflow 1%                                                                                                                                                                                                                                                                                                                                                        |
| <b>Parameters</b>  | <p><i>percent</i> — Specifies the percentage of the flow cache entries removed when the maximum number of entries is exceeded.</p> <p><b>Values</b> 1 to 50 percent</p>                                                                                                                                                                                            |

## rate

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>rate</b> <i>sample-rate</i></p> <p><b>no rate</b></p>                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Context</b>     | config>cflowd                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b> | <p>This command specifies the rate (N) at which traffic is sampled and sent for flow analysis. A packet is sampled every N packets; for example, when <b>sample-rate</b> is configured as 1, then all packets are sent to the cache. When <b>sample-rate</b> is configured as 100, then every 100th packet is sent to the cache.</p> <p>The <b>no</b> form of this command resets the sample rate to the default value.</p> |
| <b>Default</b>     | rate 1000                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Parameters</b>  | <p><i>sample-rate</i> — Specifies the rate at which traffic is sampled.</p> <p><b>Values</b> 1 to 10000</p>                                                                                                                                                                                                                                                                                                                 |

## template-retransmit

|                    |                                                                                                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <p><b>template-retransmit</b> <i>seconds</i></p> <p><b>no template-retransmit</b></p>                                                |
| <b>Context</b>     | config>cflowd                                                                                                                        |
| <b>Description</b> | This command specifies the interval for sending template definitions.                                                                |
| <b>Default</b>     | template-retransmit 600                                                                                                              |
| <b>Parameters</b>  | <p><i>seconds</i> — Specifies the value expressed in seconds before sending template definitions.</p> <p><b>Values</b> 10 to 600</p> |

## use-vrtr-if-index

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>[no] use-vrtr-if-index</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | config>cflowd                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b> | <p>This command is used to export flow data using interface indexes (ifIndex values), which can be used directly as the index into the IF-MIB tables for retrieving interface statistics. Specifically, if this command is enabled, the ingressInterface (ID=10) and egressInterface (ID= 14) fields in IP flow templates used to export the flow data to cflowd version 9 and version 10 collectors will be populated with the IF-MIB ifIndex of that interface. In addition, for version 10 templates, two fields are available in the IP flow templates to specify the virtual router ID associated with the ingress and egress interfaces.</p> <p>The <b>no</b> form of this command removes the command from the active configuration and causes cflowd to return to the default behavior of populating the ingress and egress interface ID with the global IF index IDs.</p> |
| <b>Default</b>     | no use-vrtr-if-index                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |



## 6.7 Show, Tools, and Clear Command Reference

- [Command Hierarchies](#)
- [Command Descriptions](#)

### 6.7.1 Command Hierarchies

- [Show Commands](#)
- [Tools Commands](#)
- [Clear Commands](#)

#### 6.7.1.1 Show Commands

```
show
 — cflowd
 — collector [ip-address[:port]] [detail]
 — interface [ip-int-name | ip-address]
 — status
```

#### 6.7.1.2 Tools Commands

```
tools
 — dump
 — cflowd
 — cache {all | aggregate {src-dst-proto | src-dst-proto-port}} family {ipv4 |
 ipv6}
 — packet-size [ipv4 | ipv6] [clear]
 — top-flows [ipv4 | ipv6 | mpls] [clear]
 — top-protocols [clear]
```

#### 6.7.1.3 Clear Commands

```
clear
 — cflowd
```

## 6.7.2 Command Descriptions

### 6.7.2.1 Show Commands

The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

#### collector

**Syntax** `collector [ip-addr[:port]] [detail]`

**Context** `show>cflowd`

**Description** This command displays administrative and operational status of data collector configuration.

**Parameters** *ip-addr* — Display only information about the specified collector IP address.

**Default** all collectors

*:port* — Display only information the collector on the specified UDP port.

**Default** all UDP ports

**Values** 1 to 65535

**detail** — Displays details about either all collectors or the specified collector.

**Output** **cflowd Collector Output** —The following output is an example of cflowd Collector information, and [Table 87](#) describes the output fields.

```
A:R51-CfmA# show cflowd collector

=====
Cflowd Collectors
=====
Host Address Port Version AS Type Admin Oper Sent

138.120.135.103 2055 v5 peer up up 1380 records
138.120.135.103 9555 v8 origin up up 90 records
138.120.135.103 9996 v9 - up up 0 packets
138.120.214.224 2055 v5 origin up up 1380 records

Collectors : 4
=====
```

**Table 87 Show cflowd Collector Output Fields**

| Label        | Description                                                                                  |
|--------------|----------------------------------------------------------------------------------------------|
| Host Address | The IP address of a remote cflowd collector host to receive the exported cflowd data.        |
| Port         | The UDP port number on the remote cflowd collector host to receive the exported cflowd data. |
| AS Type      | The style of AS reporting used in the exported flow data.                                    |
|              | origin<br>Reflects the endpoints of the AS path which the flow is following.                 |
|              | peer<br>Reflects the AS of the previous and next hops for the flow.                          |
| Version      | Specifies the configured version for the associated collector.                               |
| Admin        | The desired administrative state for this cflowd remote collector host.                      |
| Oper         | The current operational status of this cflowd remote collector host.                         |
| Recs Sent    | The number of cflowd records that have been transmitted to this remote collector host.       |
| Collectors   | The total number of collectors using this IP address.                                        |

**cflowd Collector Detail Output** —The following output is an example of cflowd Collector information, and [Table 88](#) describes the output fields.

### Sample Output

```
A:R51-CfmA# show cflowd collector detail
=====
Cflowd Collectors (detail)
=====
Address : 138.120.135.103
Port : 2055
Description : Test v5 Collector
Version : 5
AS Type : peer
Admin State : up
Oper State : up
Records Sent : 1260
Last Changed : 09/03/2009 17:24:04
Last Pkt Sent : 09/03/2009 18:07:10

Sent Open Errors
```

```

 42 0 0
=====
Address : 138.120.135.103
Port : 9555
Description : Test v8 Collector
Version : 8
AS Type : origin
Admin State : up
Oper State : up
Records Sent : 82
Last Changed : 09/03/2009 17:24:04
Last Pkt Sent : 09/03/2009 18:06:41

Aggregation Type Status Sent Open Errors

as-matrix Disabled 0 0 0
protocol-port Disabled 0 0 0
source-prefix Enabled 21 0 0
destination-prefix Enabled 21 0 0
source-destination-prefix Disabled 0 0 0
raw Disabled 0 0 0
=====
Address : 138.120.135.103
Port : 9996
Description : Test v9 Collector
Version : 9
Admin State : up
Oper State : up
Packets Sent : 51
Last Changed : 09/03/2009 17:24:04
Last Pkt Sent : 09/03/2009 18:07:10
Template Set : Basic

Traffic Type Template Sent Sent Open Errors

IPv4 09/03/2009 18:07:29 51 1 0
MPLS No template sent 0 0 0
IPv6 No template sent 0 0 0
=====
A:R51-CfmA#

```

**Table 88** Show cflowd Collector Detailed Output Fields

| Label       | Description                                                                                  |
|-------------|----------------------------------------------------------------------------------------------|
| Address     | The IP address of a remote cflowd collector host to receive the exported cflowd data.        |
| Port        | The UDP port number on the remote cflowd collector host to receive the exported cflowd data. |
| Description | A user-provided descriptive string for this cflowd remote collector host.                    |
| Version     | The version of the flow data sent to the collector.                                          |

**Table 88 Show cflowd Collector Detailed Output Fields (Continued)**

| Label            | Description                                                                                                                                                                          |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AS Type          | The style of AS reporting used in the exported flow data.                                                                                                                            |
|                  | origin<br>Reflects the endpoints of the AS path which the flow is following.                                                                                                         |
|                  | peer<br>Reflects the AS of the previous and next hops for the flow.                                                                                                                  |
| Admin State      | The desired administrative state for this cflowd remote collector host.                                                                                                              |
| Oper State       | The current operational status of this cflowd remote collector host.                                                                                                                 |
| Records Sent     | The number of cflowd records that have been transmitted to this remote collector host.                                                                                               |
| Last Changed     | The time when this row entry was last changed.                                                                                                                                       |
| Last Pkt Sent    | The time when the last cflowd packet was sent to this remote collector host.                                                                                                         |
| Aggregation Type | The bit mask which specifies the aggregation schemes used to aggregate multiple individual flows into an aggregated flow for export to this remote host collector.                   |
|                  | none<br>No data will be exported for this remote collector host.                                                                                                                     |
|                  | raw<br>Flow data is exported without aggregation in version 5 format.                                                                                                                |
|                  | All other aggregation types use version 8 format to export the flow data to this remote host collector.                                                                              |
| Collectors       | The total number of collectors using this IP address.                                                                                                                                |
| Sent             | The number of packets with flow data sent to the associated collector.                                                                                                               |
| Open             | This counter shows the number of partially filled packets which have some flow data but are not yet filled or have been timed out (60 seconds maximum).                              |
| Error            | This counter increments when there was an error during exporting of the collector packet. The most common reason will be a UDP unreachable destination for the configured collector. |

## interface

- Syntax** `interface [ip-addr | ip-int-name]`
- Context** `show>cflowd`
- Description** Displays the administrative and operational status of the interfaces with cflowd enabled.
- Parameters**
- ip-addr* — Display only information for the IP interface with the specified IP address.
    - Default** all interfaces with cflowd enabled.
  - ip-int-name* — Display only information for the IP interface with the specified name.
    - Default** all interfaces with cflowd enabled.
- Output** The following output is an example of cflowd interface information, and [Table 89](#) describes the output fields.

**Table 89 Show cflowd Interface Output Fields**

| Label        | Description                                                                                                                                                                                                    |
|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Interface    | Displays the physical port identifier.                                                                                                                                                                         |
| IPv4 Address | Displays the primary IPv4 address for the associated IP interface.                                                                                                                                             |
| IPv6 Address | Displays the primary IPv6 address for the associated IP interface.                                                                                                                                             |
| Router       | Displays the virtual router index (Base = 0).                                                                                                                                                                  |
| IF Index     | Displays the Global IP interface index.                                                                                                                                                                        |
| Mode         | Displays the cflowd sampling type and direction.<br>intf — Interface based sampling<br>acl — ACL based sampling<br>ingr — Ingress sampling<br>egr — Egress sampling<br>both — Both ingress and egress sampling |
| Admin        | Displays the administrative state of the interface.                                                                                                                                                            |
| Opr-IPv4     | Displays the operational state for IPv4 sampling.                                                                                                                                                              |
| Opr-IPv6     | Displays the operational state for IPv6 sampling.                                                                                                                                                              |

### Sample Output

```

B:sr-002# show cflowd interface [ip-addr | ip-int-name]
=====
Cflowd Interfaces
=====
Interface Router IF Index Mode Admin

```

| IPv4 Address     |       |       |          | Oper IPv4 |
|------------------|-------|-------|----------|-----------|
| IPv6 Address     |       |       |          | Oper IPv6 |
| -----            | ----- | ----- | -----    | -----     |
| ipv4ipv6NamedIf  | Base  | 381   | intf/ing | Up        |
| 5.5.5.5/24       |       |       |          | Up        |
| 55::55/128       |       |       |          | Up        |
| ipv4NamedIf      | 5     | 254   | acl-egr  | Up        |
| 10.10.10.10/24   |       |       |          | Up        |
| N/A              |       |       |          | Down      |
| ipv6NamedIf      | Base  | 380   | i/f-both | Up        |
| N/A              |       |       |          | Down      |
| 1234:5678::9/128 |       |       |          | Up        |
| -----            | ----- | ----- | -----    | -----     |

Interfaces : 3

B:sr-002# show cflowd interface 11.10.1.2

Cflowd Interfaces

Interface: To\_Sr1  
IP address: 11.10.1.2/24  
Admin/Oper state: Up/Up  
Sampling Mode: (ingress | egress | both)  
Total Flows seen: 1302000  
Pkts sampled (ingress/egress) : 60103/70102  
Bytes sampled (ingress/egress) : 6010300/7010200  
Active flows (ingress/egress) : 6010/7010

B:sr-002# show cflowd interface

Cflowd Interfaces

| Interface     | IP Address     | Mode      | Admin | Oper  |
|---------------|----------------|-----------|-------|-------|
| -----         | -----          | -----     | ----- | ----- |
| To_Sr1        | 1.10.1.2/24    | Interface | Up    | Up    |
| To_C2         | 1.12.1.2/24    | Interface | Up    | Up    |
| To_Cisco_7600 | 1.13.1.2/24    | Interface | Up    | Up    |
| To_E          | 1.11.1.2/24    | Interface | Up    | Up    |
| To_G2         | 150.153.1.1/24 | Interface | Up    | Up    |
| To_Sr1_Sonet  | 150.140.1.2/24 | Interface | Up    | Down  |
| Main          | 120.1.1.1/24   | Filter    | Down  | Down  |
| New           | 120.2.1.1/24   | Filter    | Up    | Up    |
| -----         | -----          | -----     | ----- | ----- |

Interfaces : 8

B:sr12-002#

## status

|                |               |
|----------------|---------------|
| <b>Syntax</b>  | <b>status</b> |
| <b>Context</b> | show>cflowd   |

**Description** This command displays basic information regarding the administrative and operational status of cflowd.

**Output** The following output is an example of cflowd status information, and [Table 90](#) describes the output fields.

### Sample Output

```
srl# show cflowd status
=====
Cflowd Status
=====
Cflowd Admin Status : Enabled
Cflowd Oper Status : Enabled
Active Timeout : 1 minutes
Inactive Timeout : 30 seconds
Template Retransmit : 60 seconds
Cache Size : 65536 entries
Overflow : 1%
Sample Rate : 1
Active Flows : 34000
Overflow events 10
Dropped Flows: 0
Pkts Rcvd : 801600
Total Pkts Dropped : 0

 Raw
Times flow created 160000
Times flow matched 224428382
Total flows flushed 150000
=====
Version Info
=====
Version Status Sent Open Errors

5 Enabled 92 0 0
8 Enabled 46 0 0
9 Enabled 56 1 0
10 Enabled 39 1 0
=====

=====
Cflowd Status
=====
Cflowd Admin Status : Enabled
Cflowd Oper Status : Enabled
Active Timeout : 1 minutes
Inactive Timeout : 30 seconds
Template Retransmit : 60 seconds
Cache Size : 65536 entries
Overflow : 1%
Sample Rate : 1
Active Flows : 34
Total Pkts Rcvd : 801600
Total Pkts Dropped : 0

=====
Version Info
=====
```



```
=====
Version Status Sent Open Errors

 5 Enabled 92 0 0
 8 Enabled 46 0 0
 9 Enabled 56 1 0
 10 Enabled 39 1 0
=====
```

**Table 90**      **Show cflowd Status Fields**

| Label               | Description                                                                                                                                                                              |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cflowd Admin Status | The desired administrative state for this cflowd remote collector host.                                                                                                                  |
| Cflowd Oper Status  | The current operational status of this cflowd remote collector host.                                                                                                                     |
| Active Timeout      | The maximum amount of time, in minutes, before an active flow will be exported. If an individual flow is active for this amount of time, the flow is exported and a new flow is created. |
| Inactive Timeout    | Inactive timeout in seconds.                                                                                                                                                             |
| Template Retransmit | The time in seconds before template definitions are sent.                                                                                                                                |
| Cache Size          | The maximum number of active flows to be maintained in the flow cache table.                                                                                                             |
| Overflow            | The percentage number of flows to be flushed when the flow cache size has been exceeded.                                                                                                 |
| Sample Rate         | The rate at which traffic is sampled and forwarded for cflowd analysis.<br>one (1)<br>All packets are analyzed.<br>1000 (default)<br>Every 1000th packet is analyzed.                    |
| Active Flows        | The current number of active flows being collected.                                                                                                                                      |
| Total Pkts Rcvd     | The total number of packets sampled and forwarded for cflowd analysis.                                                                                                                   |
| Total Pkts Dropped  | The total number of packets dropped.                                                                                                                                                     |
| Aggregation Info:   |                                                                                                                                                                                          |
| Type                | The type of data to be aggregated and to the collector.                                                                                                                                  |

**Table 90 Show cflowd Status Fields (Continued)**

| Label           | Description                                                                                                                                                                          |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Status          | enabled<br>Specifies that the aggregation type is enabled.                                                                                                                           |
|                 | disabled<br>Specifies that the aggregation type is disabled.                                                                                                                         |
| Sent            | The number of packets with flow data sent to the associated collector.                                                                                                               |
| Open            | This counter shows the number of partially filled packets which have some flow data but are not yet filled or have been timed out (60 seconds maximum).                              |
| Error           | This counter increments when there was an error during exporting of the collector packet. The most common reason will be a UDP unreachable destination for the configured collector. |
| Overflow events | The number of times the active cache overflowed.                                                                                                                                     |
| Dropped Flows   | Total number of flows dropped due to cache overflow events.                                                                                                                          |

### 6.7.2.2 Tools Commands

The following command outputs are examples only; actual displays may differ depending on supported functionality and user configuration.

#### cache

**Syntax** `cache {all | aggregate {src-dst-proto | src-dst-proto-port}} family {ipv4 | ipv6}`

**Context** `tools>dump>cflowd`

**Description** This command displays the contents of the cflowd active cache. This information can be displayed either in raw form where every flow entry is displayed or in an aggregated form.

[Table 91](#) describes the cflowd cache output fields.

**Table 91 Tools Dump cflowd Cache Output Fields**

| Label          | Description                              |
|----------------|------------------------------------------|
| Proto/Protocol | Displays the IPv4 or IPv6 protocol type. |

**Table 91 Tools Dump cflowd Cache Output Fields (Continued)**

| Label                      | Description                                                                                                                                        |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| Source Address/Src-IP      | Displays the source IP address of the flow (IPv4 or IPv6).                                                                                         |
| Destination Address/Dst-IP | Displays the destination IP address of the flow (IPv4 or IPv6).                                                                                    |
| Intf/Ingr                  | Displays the ingress interface associated with the sampled flow (only displayed with the raw (all) output).                                        |
| Intf/Egr                   | Displays the egress interface associated with the sampled flow (only displayed with the raw (all) output).                                         |
| S-Port                     | Displays the source protocol port number.                                                                                                          |
| D-Port                     | Displays the destination protocol port number.                                                                                                     |
| Pkt-Cnt                    | Displays the total number of packets sampled for the associated flow.                                                                              |
| Byte-Cnt                   | Displays the total number of bytes of traffic sampled for the associated flow.                                                                     |
| Start-Time                 | Displays the system time when the first packet was sampled for the associated flow.                                                                |
| Flags                      | Displays the IP flag value from the sampled IP flow header (only displayed with the raw (all) output).                                             |
| ToS                        | Displays the ToS byte values from the sampled IP flow header (only displayed with the raw (all) output).                                           |
| (Src) Mask                 | Displays the IP route mask for the route to the flow source IP address associated with the flow (only displayed with the raw (all) output).        |
| (Dst) Mask                 | Displays the IP route mask for the route to the flow destination IP address associated with the flow (only displayed with the raw (all) output).   |
| (Src) AS                   | Displays the ASN associated with the route to the flow source IP address associated with the flow (only displayed with the raw (all) output).      |
| (Dst) AS                   | Displays the ASN associated with the route to the flow destination IP address associated with the flow (only displayed with the raw (all) output). |
| vRtr-ID                    | Displays the Virtual Router ID associated with the reported IP flow (only displayed with the raw (all) output).                                    |

- Parameters**
- all** — Display the raw active cache flow data with no aggregation.
  - aggregate** — Display the aggregated active cache flow data.
    - src-dst-proto** — Aggregates the active flow cache based on the source and destination IP address and the IP protocol value.
    - src-dst-proto-port** — Aggregates the active flow cache based on the source and destination IP address, IP protocol value, and the source and destination port numbers.
  - family** — Specifies which IP address family flow data should be displayed.
    - ipv4** — Displays the IPv4 flow data.
    - ipv6** — Displays the IPv6 flow data.

## packet-size

- Syntax** **packet-size [ipv4 | ipv6] [clear]**
- Context** tools>dump>cflowd
- Description** This command displays packet size distribution for sampled IP traffic. Values are displays in decimal format (1.0 = 100%, .500 = 50%). Separate statistics are maintained and shown for IPv4 and IPv6 traffic.
- Output** The following output is an example of cflowd packet size information.

### Sample Output

```
SR-12# tools dump cflowd packet-size ipv4
IP packet size distribution (801600 total packets):
 1-32 64 96 128 160 192 224 256 288 320 352 384 416 448 480
 .000 .250 .000 .000 .010 .100 .500 .090 .000 .000 .000 .000 .000 .000
 512 544 576 1024 1536 2048 2560 3072 3584 4096 4608 9000
 .000 .000 .000 .050 .000 .000 .000 .000 .000 .000 .000 .000
```

## top-flows

- Syntax** **top-flows [ipv4 | ipv6 | mpls] [clear]**
- Context** tools>dump>cflowd
- Description** This command displays the top 20 (highest traffic volume) flows for IPv4, IPv6 or MPLS traffic types collected since the cflowd top-flow table was last cleared or initialized.
- Output** The following output is an example of cflowd top flow information, and [Table 92](#) describes the output fields.

### Sample Output

```

1 2 3 4 5 6 7 8
1234567890123456789012345678901234567890123456789012345678901234567890
Sr1# tools dump cflowd top-flows ipv4

Ingress i/f SrcIP Egress i/f DstIP Pr TOS Flgs Pkts
vRtr-ID S-Port Msk AS D-Port Msk AS NextHop Avg Pkt Size Active

1000 52.52.52.1 2001 123.123.123.122 0x01 55 0x10 3748
10201 0000 /8 50 0000 /8 40 202.120.130.2 220 3600
.....

1 2 3 4 5 6 7 8
1234567890123456789012345678901234567890123456789012345678901234567890
Sr1# tools dump cflowd top-flows ipv6
SrcIP (up to IPv6) Ingress i/f Src Port vRtr ID ToS
DstIP (upto IPv6) Egress i/f Dst Port Proto Flags
Nexthop (uptoIPv6) Total Pkts Avg Pkt Active(sec)
2001:0db8:85a3:0000:0000:8a2e:0370:7334 60005 10020 0 0x12
2001:0db8:85a3:0000:0000:8a2e:0280:1234 60325 20010 17 0x23
2001:0db8:85a3:0000:0000:8a2e:1234:5678 1234567890 1500 13600
.....

1 2 3 4 5 6 7 8
1234567890123456789012345678901234567890123456789012345678901234567890
Sr1# tools dump cflowd top-flows mpls
Label-1 Label-2 Label-3 Label-4 Total Pkts Avg Pkt Active(s)
SrcIP (up to IPv6) Ingress i/f Src Port ToS
DstIP (upto IPv6) Egress i/f Dst Port Proto Flags

```

**Table 92 Tools Dump cflowd Top-flows Out put Fields**

| Label       | Description                                                                                 |
|-------------|---------------------------------------------------------------------------------------------|
| Ingress     | Displays the ingress interface ID.                                                          |
| Src IP      | Displays the source IP address of the flow (IPv4 or IPv6).                                  |
| Egress      | Displays the egress interface ID.                                                           |
| Dest IP     | Displays the destination IP address of the flow (IPv4 or IPv6).                             |
| Pr<br>Proto | Displays the protocol type for flow.                                                        |
| TOS         | Displays the Type of Service/DSCP butts filed markings.                                     |
| Flgs        | Displays the protocol flag markings.                                                        |
| Pkts        | Displays the total number of packets sampled for this flow (since stats were last cleared). |
| vRtr-ID     | Displays the vRouter context the flow was sample in.                                        |

**Table 92 Tools Dump cflowd Top-flows Out put Fields (Continued)**

| Label              | Description                                                                                                                                        |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| S-Port<br>Src Port | Displays the source protocol port number.                                                                                                          |
| Msk                | Displays the route prefix length for route to source IP address.                                                                                   |
| AS                 | Displays the Autonomous Systems number for the source route (the AS is either originating AS or peer AS depending on cflowd configuration).        |
| D-Port<br>Dst Port | Displays the destination protocol port number.                                                                                                     |
| Msk                | Displays the route prefix length for route to destination IP address (Forwarding route).                                                           |
| AS                 | Displays the Autonomous Systems number for the destination route (the AS is either originating AS or peer AS depending on cflowd configuration)    |
| Nexthop            | Displays the next-hop address used to forward traffic associated with the flow.                                                                    |
| Avg pkt size       | Displays the average packet size of a sampled traffic associated with this flow (total number of packets sampled/total number of packets sampled). |
| Active             | Displays the number of seconds the flow has been active.                                                                                           |

## top-protocols

**Syntax** **top-protocols**

**Context** tools>dump>cflowd [clear]

**Description** This command displays the summary information for the top 20 protocol traffic seen in the cflowd cache. All statistics are calculated based on the data collected since the last clearing of the cflowd stats with clear keyword for this command.

If the clear optional keyword is given, then the top-flows are displayed, and then this cache is cleared.

**Output** The following output is an example of Cflowd top protocol traffic information, and [Table 93](#) describes the output fields.

### Sample Output

```
SR# tools dump cflowd top-protocols
```

The top 20 IPv4 protocols seen by cflowd are:

Current Time: 08/29/2011 15:36:15

Last Cleared Time: 08/29/2011 15:35:08

| Protocol ID<br>----- | Total<br>Flows | Flows<br>/Sec | Packets<br>/Flow | Bytes<br>/Pkt | Packets<br>/Sec | Duration<br>/Flow | % Total<br>Bandwidth |
|----------------------|----------------|---------------|------------------|---------------|-----------------|-------------------|----------------------|
| UDP                  | 2              | 0             | 6                | 100           | 0               | 6                 | 75%                  |
| prl                  | 1              | 0             | 6                | 64            | 0               | 6                 | 24%                  |
| TOTALS               | 3              | 0             | 6                | 88            | 0               | 6                 | 100%                 |

**Table 93 Tools Dump cflowd Top-protocols Fields**

| Label               | Description                                                                                                                                                                                       |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Protocol ID         | Displays the IPv4 or IPv6 protocol type.<br>This will either print the well-known protocol name or the decimal protocol number.                                                                   |
| Total Flows         | Displays the total number of flows recorded since the last clearing of cflowd statistics with this protocol type.                                                                                 |
| Flows/Sec           | Displays the average number of flows detected for the associated protocol type.<br>(Total flows/number of seconds since last clear)                                                               |
| Packets/Flow        | Displays the average number of packets per flow.<br>(Total number of packets/total flows)                                                                                                         |
| Bytes/Pkts          | Displays the average number of bytes per packet for the associated protocol type.<br>(Total number of bytes for the associated protocol/total number of packets seen for the associated protocol) |
| Packets/Sec         | Displays the average number of packets seen for the associated protocol type.<br>(Number of packets/time since last clear)                                                                        |
| Duration/Flow       | Displays the average lifetime of a flow for the associated protocol type.<br>(Number of seconds since last clear/total flows)                                                                     |
| Bandwidth Total (%) | Displays the percentage of bandwidth consumed by the associated protocol type.<br>(Total protocol bytes/total bytes of all flows)                                                                 |

### 6.7.2.3 Clear Commands

#### cflowd

|                    |                                                                                                                                                                                                                                                                                                                    |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>      | <b>cflowd</b>                                                                                                                                                                                                                                                                                                      |
| <b>Context</b>     | clear                                                                                                                                                                                                                                                                                                              |
| <b>Description</b> | Clears the raw and aggregation flow caches which are sending flow data to the configured collectors. This action will trigger all the flows to be discarded. The cache restarts flow data collection from a fresh state. This command also clears global stats collector stats listed in the cflowd show commands. |



## 7 Standards and Protocol Support



**Note:** The information presented is subject to change without notice.

Nokia assumes no responsibility for inaccuracies contained herein.

### Access Node Control Protocol (ANCP)

draft-ietf-ancp-protocol-02, *Protocol for Access Node Control Mechanism in Broadband Networks*

RFC 5851, *Framework and Requirements for an Access Node Control Mechanism in Broadband Multi-Service Networks*

### Application Assurance (AA)

3GPP Release 12 (ADC rules over Gx interfaces)

RFC 3507, *Internet Content Adaptation Protocol (ICAP)*

### Asynchronous Transfer Mode (ATM)

AF-ILMI-0065.000, *Integrated Local Management Interface (ILMI) Version 4.0*

AF-PHY-0086.001, *Inverse Multiplexing for ATM (IMA) Specification Version 1.1*

AF-TM-0121.000, *Traffic Management Specification Version 4.1*

AF-TM-0150.00, *Addendum to Traffic Management v4.1 optional minimum desired cell rate indication for UBR*

GR-1113-CORE, *Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1*

GR-1248-CORE, *Generic Requirements for Operations of ATM Network Elements (NEs), Issue 3*

ITU-T I.432.1, *B-ISDN user-network interface - Physical layer specification: General characteristics (02/99)*

ITU-T I.610, *B-ISDN operation and maintenance principles and functions (11/95)*

RFC 1626, *Default IP MTU for use over ATM AAL5*

RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*

### Bidirectional Forwarding Detection (BFD)

RFC 5880, *Bidirectional Forwarding Detection (BFD)*

RFC 5881, *Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop)*

RFC 5883, *Bidirectional Forwarding Detection (BFD) for Multihop Paths*

RFC 7130, *Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces*

## **Border Gateway Protocol (BGP)**

draft-hares-idr-update-attr-low-bits-fix-01, *Update Attribute Flag Low Bits Clarification*

draft-ietf-idr-add-paths-guidelines-08, *Best Practices for Advertisement of Multiple Paths in IBGP*

draft-ietf-idr-best-external-03, *Advertisement of the best external route in BGP*

draft-ietf-idr-bgp-flowspec-oid-03, *Revised Validation Procedure for BGP Flow Specifications*

draft-ietf-idr-bgp-gr-notification-01, *Notification Message support for BGP Graceful Restart*

draft-ietf-idr-bgp-optimal-route-reflection-10, *BGP Optimal Route Reflection (BGP-ORR)*

draft-ietf-idr-error-handling-03, *Revised Error Handling for BGP UPDATE Messages*

draft-ietf-idr-flowspec-interfaceset-03, *Applying BGP flowspec rules on a specific interface set*

draft-ietf-idr-link-bandwidth-03, *BGP Link Bandwidth Extended Community*

draft-ietf-sidr-origin-validation-signaling-04, *BGP Prefix Origin Validation State Extended Community*

draft-uttaro-idr-bgp-persistence-03, *Support for Long-lived BGP Graceful Restart*

RFC 1772, *Application of the Border Gateway Protocol in the Internet*

RFC 1997, *BGP Communities Attribute*

RFC 2385, *Protection of BGP Sessions via the TCP MD5 Signature Option*

RFC 2439, *BGP Route Flap Damping*

RFC 2545, *Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing*

RFC 2858, *Multiprotocol Extensions for BGP-4*

RFC 2918, *Route Refresh Capability for BGP-4*

RFC 3107, *Carrying Label Information in BGP-4*

RFC 3392, *Capabilities Advertisement with BGP-4*

RFC 4271, *A Border Gateway Protocol 4 (BGP-4)*

RFC 4360, *BGP Extended Communities Attribute*

RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*

RFC 4456, *BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)*

RFC 4486, *Subcodes for BGP Cease Notification Message*

RFC 4659, *BGP/MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*

RFC 4684, *Constrained Route Distribution for Border Gateway Protocol/ MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)*

RFC 4724, *Graceful Restart Mechanism for BGP (helper mode)*

RFC 4760, *Multiprotocol Extensions for BGP-4*

RFC 4798, *Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)*

RFC 4893, *BGP Support for Four-octet AS Number Space*

RFC 5004, *Avoid BGP Best Path Transitions from One External to Another*

RFC 5065, *Autonomous System Confederations for BGP*

RFC 5291, *Outbound Route Filtering Capability for BGP-4*

RFC 5396, *Textual Representation of Autonomous System (AS) Numbers (asplain)*

RFC 5575, *Dissemination of Flow Specification Rules*

RFC 5668, *4-Octet AS Specific BGP Extended Community*

RFC 6810, *The Resource Public Key Infrastructure (RPKI) to Router Protocol*

RFC 6811, *Prefix Origin Validation*

RFC 6996, *Autonomous System (AS) Reservation for Private Use*

RFC 7311, *The Accumulated IGP Metric Attribute for BGP*

RFC 7607, *Codification of AS 0 Processing*

RFC 7674, *Clarification of the Flowspec Redirect Extended Community*

RFC 7752, *North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP*

RFC 7911, *Advertisement of Multiple Paths in BGP*

## **Circuit Emulation**

RFC 4553, *Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)*

RFC 5086, *Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)*

RFC 5287, *Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks*

## **Ethernet**

IEEE 802.1AB, *Station and Media Access Control Connectivity Discovery*

IEEE 802.1ad, *Provider Bridges*

IEEE 802.1ag, *Connectivity Fault Management*

IEEE 802.1ah, *Provider Backbone Bridges*

---

IEEE 802.1ak, *Multiple Registration Protocol*  
IEEE 802.1aq, *Shortest Path Bridging*  
IEEE 802.1ax, *Link Aggregation*  
IEEE 802.1D, *MAC Bridges*  
IEEE 802.1p, *Traffic Class Expediting*  
IEEE 802.1Q, *Virtual LANs*  
IEEE 802.1s, *Multiple Spanning Trees*  
IEEE 802.1w, *Rapid Reconfiguration of Spanning Tree*  
IEEE 802.1X, *Port Based Network Access Control*  
IEEE 802.3ab, *1000BASE-T*  
IEEE 802.3ac, *VLAN Tag*  
IEEE 802.3ad, *Link Aggregation*  
IEEE 802.3ae, *10 Gb/s Ethernet*  
IEEE 802.3ah, *Ethernet in the First Mile*  
IEEE 802.3ba, *40 Gb/s and 100 Gb/s Ethernet*  
IEEE 802.3i, *Ethernet*  
IEEE 802.3u, *Fast Ethernet*  
IEEE 802.3x, *Ethernet Flow Control*  
IEEE 802.3z, *Gigabit Ethernet*  
ITU-T G.8031/Y.1342, *Ethernet Linear Protection Switching*  
ITU-T G.8032/Y.1344, *Ethernet Ring Protection Switching*  
ITU-T Y.1731, *OAM functions and mechanisms for Ethernet based networks*

## **Ethernet VPN (EVPN)**

draft-ietf-bess-evpn-ac-df-01, *AC-Influenced Designated Forwarder Election for EVPN*  
draft-ietf-bess-evpn-etree-11, *E-TREE Support in EVPN & PBB-EVPN*  
draft-ietf-bess-evpn-overlay-04, *A Network Virtualization Overlay Solution using EVPN*  
draft-ietf-bess-evpn-prefix-advertisement-02, *IP Prefix Advertisement in EVPN*  
draft-ietf-bess-evpn-proxy-arp-nd-02, *Operational Aspects of Proxy-ARP/ND in EVPN Networks*  
draft-ietf-bess-evpn-vpls-seamless-integ-00, *(PBB-)EVPN Seamless Integration with (PBB-)VPLS*  
draft-ietf-bess-evpn-vpws-14, *Virtual Private Wire Service support in Ethernet VPN*  
draft-rabadan-bess-evpn-pref-df-02, *Preference-based EVPN DF Election*  
draft-snr-bess-pbb-evpn-isid-cmacflush-01, *PBB-EVPN ISID-based CMAC-Flush*

RFC 7432, *BGP MPLS-Based Ethernet VPN*

RFC 7623, *Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)*

## **Frame Relay**

ANSI T1.617 Annex D, *DSS1 - Signalling Specification For Frame Relay Bearer Service*

FRF.1.2, *PVC User-to-Network Interface (UNI) Implementation Agreement*

FRF.12, *Frame Relay Fragmentation Implementation Agreement*

FRF.16.1, *Multilink Frame Relay UNI/NNI Implementation Agreement*

FRF.5, *Frame Relay/ATM PVC Network Interworking Implementation*

FRF2.2, *PVC Network-to-Network Interface (NNI) Implementation Agreement*

ITU-T Q.933 Annex A, *Additional procedures for Permanent Virtual Connection (PVC) status management*

## **Generalized Multiprotocol Label Switching (GMPLS)**

draft-ietf-ccamp-rsvp-te-srlg-collect-04, *RSVP-TE Extensions for Collecting SRLG Information*

RFC 3471, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description*

RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions*

RFC 4204, *Link Management Protocol (LMP)*

RFC 4208, *Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model*

RFC 4872, *RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery*

RFC 5063, *Extensions to GMPLS Resource Reservation Protocol (RSVP) Graceful Restart (helper mode)*

## **Intermediate System to Intermediate System (IS-IS)**

draft-ginsberg-isis-mi-bis-01, *IS-IS Multi-Instance (single topology)*

draft-ietf-isis-mi-02, *IS-IS Multi-Instance*

draft-kaplan-isis-ext-eth-02, *Extended Ethernet Frame Size Support*

ISO/IEC 10589:2002, Second Edition, Nov. 2002, *Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)*

---

RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*  
RFC 2973, *IS-IS Mesh Groups*  
RFC 3359, *Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System*  
RFC 3719, *Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)*  
RFC 3787, *Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS)*  
RFC 4971, *Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information*  
RFC 5120, *M-ISIS: Multi Topology (MT) Routing in IS-IS*  
RFC 5130, *A Policy Control Mechanism in IS-IS Using Administrative Tags*  
RFC 5301, *Dynamic Hostname Exchange Mechanism for IS-IS*  
RFC 5302, *Domain-wide Prefix Distribution with Two-Level IS-IS*  
RFC 5303, *Three-Way Handshake for IS-IS Point-to-Point Adjacencies*  
RFC 5304, *IS-IS Cryptographic Authentication*  
RFC 5305, *IS-IS Extensions for Traffic Engineering TE*  
RFC 5306, *Restart Signaling for IS-IS (helper mode)*  
RFC 5307, *IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*  
RFC 5308, *Routing IPv6 with IS-IS*  
RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*  
RFC 5310, *IS-IS Generic Cryptographic Authentication*  
RFC 6213, *IS-IS BFD-Enabled TLV*  
RFC 6232, *Purge Originator Identification TLV for IS-IS*  
RFC 6233, *IS-IS Registry Extension for Purges*  
RFC 6329, *IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging*  
RFC 7775, *IS-IS Route Preference for Extended IP and IPv6 Reachability*  
RFC 7794, *IS-IS Prefix Attributes for Extended IPv4 and IPv6 Reachability*

## **Internet Protocol (IP) — Fast Reroute**

draft-ietf-rtgwg-lfa-manageability-08, *Operational management of Loop Free Alternates*  
RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*  
RFC 7431, *Multicast-Only Fast Reroute*  
RFC 7490, *Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)*

---

## Internet Protocol (IP) — General

draft-grant-tacacs-02, *The TACACS+ Protocol*  
RFC 768, *User Datagram Protocol*  
RFC 793, *Transmission Control Protocol*  
RFC 854, *Telnet Protocol Specifications*  
RFC 1350, *The TFTP Protocol (revision 2)*  
RFC 2347, *TFTP Option Extension*  
RFC 2348, *TFTP Blocksize Option*  
RFC 2349, *TFTP Timeout Interval and Transfer Size Options*  
RFC 2428, *FTP Extensions for IPv6 and NATs*  
RFC 2784, *Generic Routing Encapsulation (GRE)*  
RFC 4250, *The Secure Shell (SSH) Protocol Assigned Numbers*  
RFC 4251, *The Secure Shell (SSH) Protocol Architecture*  
RFC 4252, *The Secure Shell (SSH) Authentication Protocol (publickey, password)*  
RFC 4253, *The Secure Shell (SSH) Transport Layer Protocol*  
RFC 4254, *The Secure Shell (SSH) Connection Protocol*  
RFC 4632, *Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan*  
RFC 5082, *The Generalized TTL Security Mechanism (GTSM)*  
RFC 5656, *Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer (ECDSA)*  
RFC 6398, *IP Router Alert Considerations and Usage (MLD)*  
RFC 6528, *Defending against Sequence Number Attacks*

## Internet Protocol (IP) — Multicast

cisco-ipmulticast/pim-autorp-spec01, *Auto-RP: Automatic discovery of Group-to-RP mappings for IP multicast (version 1)*  
draft-dolganow-bess-mvpn-expl-track-01, *Explicit Tracking with Wild Card Routes in Multicast VPN*  
draft-ietf-idmr-traceroute-ipm-07, *A "traceroute" facility for IP Multicast*  
draft-ietf-l2vpn-vpls-pim-snooping-07, *Protocol Independent Multicast (PIM) over Virtual Private LAN Service (VPLS)*  
RFC 1112, *Host Extensions for IP Multicasting*  
RFC 2236, *Internet Group Management Protocol, Version 2*  
RFC 2365, *Administratively Scoped IP Multicast*  
RFC 2375, *IPv6 Multicast Address Assignments*  
RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*

---

RFC 3306, *Unicast-Prefix-based IPv6 Multicast Addresses*  
RFC 3376, *Internet Group Management Protocol, Version 3*  
RFC 3446, *Anycast Rendezvous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*  
RFC 3590, *Source Address Selection for the Multicast Listener Discovery (MLD) Protocol*  
RFC 3618, *Multicast Source Discovery Protocol (MSDP)*  
RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*  
RFC 3956, *Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address*  
RFC 3973, *Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised) (auto-RP groups)*  
RFC 4541, *Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches*  
RFC 4601, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*  
RFC 4604, *Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast*  
RFC 4607, *Source-Specific Multicast for IP*  
RFC 4608, *Source-Specific Protocol Independent Multicast in 232/8*  
RFC 4610, *Anycast-RP Using Protocol Independent Multicast (PIM)*  
RFC 4611, *Multicast Source Discovery Protocol (MSDP) Deployment Scenarios*  
RFC 5059, *Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)*  
RFC 5186, *Internet Group Management Protocol Version 3 (IGMPv3) / Multicast Listener Discovery Version 2 (MLDv2) and Multicast Routing Protocol Interaction*  
RFC 5384, *The Protocol Independent Multicast (PIM) Join Attribute Format*  
RFC 5496, *The Reverse Path Forwarding (RPF) Vector TLV*  
RFC 6037, *Cisco Systems' Solution for Multicast in MPLS/BGP IP VPNs*  
RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root*  
RFC 6513, *Multicast in MPLS/BGP IP VPNs*  
RFC 6514, *BGP Encodings and Procedures for Multicast in MPLS/IP VPNs*  
RFC 6515, *IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs*  
RFC 6516, *IPv6 Multicast VPN (MVPN) Support Using PIM Control Plane and Selective Provider Multicast Service Interface (S-PMSI) Join Messages*  
RFC 6625, *Wildcards in Multicast VPN Auto-Discover Routes*



RFC 6826, *Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path*

RFC 7246, *Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF) Table Context*

RFC 7385, *IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points*

RFC 7716, *Global Table Multicast with BGP Multicast VPN (BGP-MVPN) Procedures*

## **Internet Protocol (IP) — Version 4**

RFC 791, *Internet Protocol*

RFC 792, *Internet Control Message Protocol*

RFC 826, *An Ethernet Address Resolution Protocol*

RFC 951, *Bootstrap Protocol (BOOTP)*

RFC 1034, *Domain Names - Concepts and Facilities*

RFC 1035, *Domain Names - Implementation and Specification*

RFC 1519, *Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*

RFC 1534, *Interoperation between DHCP and BOOTP*

RFC 1542, *Clarifications and Extensions for the Bootstrap Protocol*

RFC 1812, *Requirements for IPv4 Routers*

RFC 1918, *Address Allocation for Private Internets*

RFC 2003, *IP Encapsulation within IP*

RFC 2131, *Dynamic Host Configuration Protocol*

RFC 2132, *DHCP Options and BOOTP Vendor Extensions*

RFC 2401, *Security Architecture for Internet Protocol*

RFC 3021, *Using 31-Bit Prefixes on IPv4 Point-to-Point Links*

RFC 3046, *DHCP Relay Agent Information Option (Option 82)*

RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*

RFC 4884, *Extended ICMP to Support Multi-Part Messages (ICMPv4 and ICMPv6 Time Exceeded)*

## **Internet Protocol (IP) — Version 6**

RFC 1981, *Path MTU Discovery for IP version 6*

RFC 2460, *Internet Protocol, Version 6 (IPv6) Specification*

RFC 2464, *Transmission of IPv6 Packets over Ethernet Networks*

RFC 2473, *Generic Packet Tunneling in IPv6 Specification*

---

RFC 2529, *Transmission of IPv6 over IPv4 Domains without Explicit Tunnels*  
RFC 3122, *Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification*  
RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*  
RFC 3587, *IPv6 Global Unicast Address Format*  
RFC 3596, *DNS Extensions to Support IP version 6*  
RFC 3633, *IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6*  
RFC 3646, *DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*  
RFC 3736, *Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6*  
RFC 3971, *SEcure Neighbor Discovery (SEND)*  
RFC 3972, *Cryptographically Generated Addresses (CGA)*  
RFC 4007, *IPv6 Scoped Address Architecture*  
RFC 4193, *Unique Local IPv6 Unicast Addresses*  
RFC 4291, *Internet Protocol Version 6 (IPv6) Addressing Architecture*  
RFC 4443, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*  
RFC 4861, *Neighbor Discovery for IP version 6 (IPv6)*  
RFC 4862, *IPv6 Stateless Address Autoconfiguration (router functions)*  
RFC 4890, *Recommendations for Filtering ICMPv6 Messages in Firewalls*  
RFC 4941, *Privacy Extensions for Stateless Address Autoconfiguration in IPv6*  
RFC 5007, *DHCPv6 Leasequery*  
RFC 5095, *Deprecation of Type 0 Routing Headers in IPv6*  
RFC 5798, *Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6 (IPv6)*  
RFC 5952, *A Recommendation for IPv6 Address Text Representation*  
RFC 6092 *Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service (Internet Control and Management, Upper-Layer Transport Protocols, UDP Filters, IPsec and Internet Key Exchange (IKE), TCP Filters)*  
RFC 6106, *IPv6 Router Advertisement Options for DNS Configuration*  
RFC 6164, *Using 127-Bit IPv6 Prefixes on Inter-Router Links*  
RFC 8021, *Generation of IPv6 Atomic Fragments Considered Harmful*

## **Internet Protocol Security (IPsec)**

draft-ietf-ipsec-isakmp-mode-cfg-05, *The ISAKMP Configuration Method*

---

draft-ietf-ipsec-isakmp-xauth-06, *Extended Authentication within ISAKMP/Oakley (XAUTH)*

RFC 2401, *Security Architecture for the Internet Protocol*

RFC 2403, *The Use of HMAC-MD5-96 within ESP and AH*

RFC 2404, *The Use of HMAC-SHA-1-96 within ESP and AH*

RFC 2405, *The ESP DES-CBC Cipher Algorithm With Explicit IV*

RFC 2406, *IP Encapsulating Security Payload (ESP)*

RFC 2407, *IPsec Domain of Interpretation for ISAKMP (IPsec DoI)*

RFC 2408, *Internet Security Association and Key Management Protocol (ISAKMP)*

RFC 2409, *The Internet Key Exchange (IKE)*

RFC 2410, *The NULL Encryption Algorithm and Its Use With IPsec*

RFC 3526, *More Modular Exponential (MODP) Diffie-Hellman group for Internet Key Exchange (IKE)*

RFC 3566, *The AES-XCBC-MAC-96 Algorithm and Its Use With IPsec*

RFC 3602, *The AES-CBC Cipher Algorithm and Its Use with IPsec*

RFC 3706, *A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers*

RFC 3947, *Negotiation of NAT-Traversal in the IKE*

RFC 3948, *UDP Encapsulation of IPsec ESP Packets*

RFC 4210, *Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)*

RFC 4211, *Internet X.509 Public Key Infrastructure Certificate Request Message Format (CRMF)*

RFC 4301, *Security Architecture for the Internet Protocol*

RFC 4303, *IP Encapsulating Security Payload*

RFC 4307, *Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)*

RFC 4308, *Cryptographic Suites for IPsec*

RFC 4434, *The AES-XCBC-PRF-128 Algorithm for the Internet Key Exchange Protocol (IKE)*

RFC 4868, *Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPSec*

RFC 4945, *The Internet IP Security PKI Profile of IKEv1/ISAKMP, IKEv2 and PKIX*

RFC 5019, *The Lightweight Online Certificate Status Protocol (OCSP) Profile for High-Volume Environments*

RFC 5280, *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*

RFC 5998, *An Extension for EAP-Only Authentication in IKEv2*

---

RFC 6712, *Internet X.509 Public Key Infrastructure -- HTTP Transfer for the Certificate Management Protocol (CMP)*

RFC 6960, *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP*

RFC 7296, *Internet Key Exchange Protocol Version 2 (IKEv2)*

RFC 7321, *Cryptographic Algorithm Implementation Requirements and Usage Guidance for Encapsulating Security Payload (ESP) and Authentication Header (AH)*

RFC 7383, *Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation*

RFC 7383, *Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation*

RFC 7468, *Textual Encodings of PKIX, PKCS, and CMS Structures*

## **Label Distribution Protocol (LDP)**

draft-ietf-mpls-ldp-ip-pw-capability-09, *Controlling State Advertisements Of Non-negotiated LDP Applications*

draft-pdutta-mpls-ldp-adj-capability-00, *LDP Adjacency Capabilities*

draft-pdutta-mpls-ldp-v2-00, *LDP Version 2*

draft-pdutta-mpls-mldp-up-redundancy-00, *Upstream LSR Redundancy for Multipoint LDP Tunnels*

draft-pdutta-mpls-multi-ldp-instance-00, *Multiple LDP Instances*

draft-pdutta-mpls-tldp-hello-reduce-04, *Targeted LDP Hello Reduction*

RFC 3037, *LDP Applicability*

RFC 3478, *Graceful Restart Mechanism for Label Distribution Protocol (helper mode)*

RFC 5036, *LDP Specification*

RFC 5283, *LDP Extension for Inter-Area Label Switched Paths (LSPs)*

RFC 5443, *LDP IGP Synchronization*

RFC 5561, *LDP Capabilities*

RFC 5919, *Signaling LDP Label Advertisement Completion*

RFC 6388, *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*

RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root*

RFC 6826, *Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*

RFC 7032, *LDP Downstream-on-Demand in Seamless MPLS*

RFC 7552, *Updates to LDP for IPv6*

## Layer Two Tunneling Protocol (L2TP) Network Server (LNS)

draft-mammoliti-l2tp-accessline-avp-04, *Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extensions*

RFC 2661, *Layer Two Tunneling Protocol "L2TP"*

RFC 2809, *Implementation of L2TP Compulsory Tunneling via RADIUS*

RFC 3438, *Layer Two Tunneling Protocol (L2TP) Internet Assigned Numbers: Internet Assigned Numbers Authority (IANA) Considerations Update*

RFC 3931, *Layer Two Tunneling Protocol - Version 3 (L2TPv3)*

RFC 4719, *Transport of Ethernet Frames over Layer 2 Tunneling Protocol Version 3 (L2TPv3)*

RFC 4951, *Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP) "failover"*

## Management

draft-ietf-snmpv3-update-mib-05, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*

draft-ietf-isis-wg-mib-06, *Management Information Base for Intermediate System to Intermediate System (IS-IS)*

draft-ietf-mboned-msdp-mib-01, *Multicast Source Discovery protocol MIB*

draft-ietf-mpls-ldp-mib-07, *Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP)*

draft-ietf-mpls-lsr-mib-06, *Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIv2*

draft-ietf-mpls-te-mib-04, *Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base*

draft-ietf-ospf-mib-update-08, *OSPF Version 2 Management Information Base*

draft-ietf-vrrp-unified-mib-06, *Definitions of Managed Objects for the VRRP over IPv4 and IPv6 (IPv6)*

ianaaddressfamilynumbers-mib, *IANA-ADDRESS-FAMILY-NUMBERS-MIB*

ianagmplstc-mib, *IANA-GMPLS-TC-MIB*

ianaiftype-mib, *IANAifType-MIB*

ianaiprouteprotocol-mib, *IANA-RTPROTO-MIB*

IEEE8021-CFM-MIB, *IEEE P802.1ag(TM) CFM MIB*

IEEE8021-PAE-MIB, *IEEE 802.1X MIB*

IEEE8023-LAG-MIB, *IEEE 802.3ad MIB*

LLDP-MIB, *IEEE P802.1AB(TM) LLDP MIB*

RFC 1157, *A Simple Network Management Protocol (SNMP)*

RFC 1212, *Concise MIB Definitions*

---

RFC 1213, *Management Information Base for Network Management of TCP/IP-based Internets: MIB-II*

RFC 1215, *A Convention for Defining Traps for use with the SNMP*

RFC 1724, *RIP Version 2 MIB Extension*

RFC 1901, *Introduction to Community-based SNMPv2*

RFC 2021, *Remote Network Monitoring Management Information Base Version 2 using SMIv2*

RFC 2115, *Management Information Base for Frame Relay DTEs Using SMIv2*

RFC 2206, *RSVP Management Information Base using SMIv2*

RFC 2213, *Integrated Services Management Information Base using SMIv2*

RFC 2494, *Definitions of Managed Objects for the DS0 and DS0 Bundle Interface Type*

RFC 2514, *Definitions of Textual Conventions and OBJECT-IDENTITIES for ATM Management*

RFC 2515, *Definitions of Managed Objects for ATM Management*

RFC 2570, *SNMP Version 3 Framework*

RFC 2571, *An Architecture for Describing SNMP Management Frameworks*

RFC 2572, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)*

RFC 2573, *SNMP Applications*

RFC 2574, *User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)*

RFC 2575, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*

RFC 2578, *Structure of Management Information Version 2 (SMIv2)*

RFC 2579, *Textual Conventions for SMIv2*

RFC 2580, *Conformance Statements for SMIv2*

RFC 2787, *Definitions of Managed Objects for the Virtual Router Redundancy Protocol*

RFC 2819, *Remote Network Monitoring Management Information Base*

RFC 2856, *Textual Conventions for Additional High Capacity Data Types*

RFC 2863, *The Interfaces Group MIB*

RFC 2864, *The Inverted Stack Table Extension to the Interfaces Group MIB*

RFC 2933, *Internet Group Management Protocol MIB*

RFC 3014, *Notification Log MIB*

RFC 3164, *The BSD syslog Protocol*

RFC 3165, *Definitions of Managed Objects for the Delegation of Management Scripts*

RFC 3231, *Definitions of Managed Objects for Scheduling Management Operations*

---

RFC 3273, *Remote Network Monitoring Management Information Base for High Capacity Networks*

RFC 3416, *Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)*

RFC 3417, *Transport Mappings for the Simple Network Management Protocol (SNMP) (SNMP over UDP over IPv4)*

RFC 3419, *Textual Conventions for Transport Addresses*

RFC 3498, *Definitions of Managed Objects for Synchronous Optical Network (SONET) Linear Automatic Protection Switching (APS) Architectures*

RFC 3584, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework*

RFC 3592, *Definitions of Managed Objects for the Synchronous Optical Network/ Synchronous Digital Hierarchy (SONET/SDH) Interface Type*

RFC 3593, *Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals*

RFC 3635, *Definitions of Managed Objects for the Ethernet-like Interface Types*

RFC 3637, *Definitions of Managed Objects for the Ethernet WAN Interface Sublayer*

RFC 3826, *The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model*

RFC 3877, *Alarm Management Information Base (MIB)*

RFC 3895, *Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types*

RFC 3896, *Definitions of Managed Objects for the DS3/E3 Interface Type*

RFC 4001, *Textual Conventions for Internet Network Addresses*

RFC 4022, *Management Information Base for the Transmission Control Protocol (TCP)*

RFC 4113, *Management Information Base for the User Datagram Protocol (UDP)*

RFC 4220, *Traffic Engineering Link Management Information Base*

RFC 4273, *Definitions of Managed Objects for BGP-4*

RFC 4292, *IP Forwarding Table MIB*

RFC 4293, *Management Information Base for the Internet Protocol (IP)*

RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*

RFC 4511, *Lightweight Directory Access Protocol (LDAP): The Protocol*

RFC 4513, *Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms (TLS)*

RFC 4631, *Link Management Protocol (LMP) Management Information Base (MIB)*

RFC 4878, *Definitions and Managed Objects for Operations, Administration, and Maintenance (OAM) Functions on Ethernet-Like Interfaces*

RFC 5101, *Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information*  
RFC 5102, *Information Model for IP Flow Information Export*  
RFC 5246, *The Transport Layer Security (TLS) Protocol Version 1.2* (TLS client, RSA public key)  
RFC 6424, *Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels*  
RFC 6425, *Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping*  
RFC 7420, *Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module*  
SFLOW-MIB, *sFlow MIB Version 1.3 (Draft 5)*

## **Multiprotocol Label Switching — Transport Profile (MPLS-TP)**

RFC 5586, *MPLS Generic Associated Channel*  
RFC 5921, *A Framework for MPLS in Transport Networks*  
RFC 5960, *MPLS Transport Profile Data Plane Architecture*  
RFC 6370, *MPLS Transport Profile (MPLS-TP) Identifiers*  
RFC 6378, *MPLS Transport Profile (MPLS-TP) Linear Protection*  
RFC 6426, *MPLS On-Demand Connectivity and Route Tracing*  
RFC 6427, *MPLS Fault Management Operations, Administration, and Maintenance (OAM)*  
RFC 6428, *Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile*  
RFC 6478, *Pseudowire Status for Static Pseudowires*  
RFC 7213, *MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing*

## **Multiprotocol Label Switching (MPLS)**

RFC 3031, *Multiprotocol Label Switching Architecture*  
RFC 3032, *MPLS Label Stack Encoding*  
RFC 3443, *Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks*  
RFC 4023, *Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)*  
RFC 4182, *Removing a Restriction on the use of MPLS Explicit NULL*  
RFC 5332, *MPLS Multicast Encapsulations*  
RFC 5884, *Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)*  
RFC 6790, *The Use of Entropy Labels in MPLS Forwarding*



RFC 7510, *Encapsulating MPLS in UDP*

## **Network Address Translation (NAT)**

draft-ietf-behave-address-format-10, *IPv6 Addressing of IPv4/IPv6 Translators*

draft-ietf-behave-v6v4-xlate-23, *IP/ICMP Translation Algorithm*

draft-miles-behave-l2nat-00, *Layer2-Aware NAT*

draft-nishitani-cgn-02, *Common Functions of Large Scale NAT (LSN)*

RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP*

RFC 5382, *NAT Behavioral Requirements for TCP*

RFC 5508, *NAT Behavioral Requirements for ICMP*

RFC 6146, *Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers*

RFC 6333, *Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion*

RFC 6334, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite*

RFC 6887, *Port Control Protocol (PCP)*

RFC 6888, *Common Requirements For Carrier-Grade NATs (CGNs)*

RFC 7915, *IP/ICMP Translation Algorithm*

## **Network Configuration Protocol (NETCONF)**

RFC 6020, *YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)*

RFC 6241, *Network Configuration Protocol (NETCONF)*

RFC 6242, *Using the NETCONF Protocol over Secure Shell (SSH)*

RFC 6243, *With-defaults Capability for NETCONF*

## **Open Shortest Path First (OSPF)**

draft-ietf-ospf-ospfv3-lsa-extend-13, *OSPFv3 LSA Extendibility*

RFC 1586, *Guidelines for Running OSPF Over Frame Relay Networks*

RFC 1765, *OSPF Database Overflow*

RFC 2328, *OSPF Version 2*

RFC 3101, *The OSPF Not-So-Stubby Area (NSSA) Option*

RFC 3509, *Alternative Implementations of OSPF Area Border Routers*

RFC 3623, *Graceful OSPF Restart Graceful OSPF Restart (helper mode)*

RFC 3630, *Traffic Engineering (TE) Extensions to OSPF Version 2*

---

RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*

RFC 4222, *Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance*

RFC 4552, *Authentication/Confidentiality for OSPFv3*

RFC 4576, *Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)*

RFC 4577, *OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*

RFC 5185, *OSPF Multi-Area Adjacency*

RFC 5187, *OSPFv3 Graceful Restart (helper mode)*

RFC 5243, *OSPF Database Exchange Summary List Optimization*

RFC 5250, *The OSPF Opaque LSA Option*

RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*

RFC 5340, *OSPF for IPv6*

RFC 5709, *OSPFv2 HMAC-SHA Cryptographic Authentication*

RFC 5838, *Support of Address Families in OSPFv3*

RFC 6987, *OSPF Stub Router Advertisement*

RFC 7684, *OSPFv2 Prefix/Link Attribute Advertisement*

RFC 7770, *Extensions to OSPF for Advertising Optional Router Capabilities*

## OpenConfig

gnmi.proto, *gRPC Network Management Interface (gNMI)*, version 0.3.1 (Subscribe RPC)

## OpenFlow

ONF *OpenFlow Switch Specification Version 1.3.1* (OpenFlow-hybrid switches)

## Path Computation Element Protocol (PCEP)

draft-alvarez-pce-path-profiles-04, *PCE Path Profiles*

draft-ietf-pce-segment-routing-08, *PCEP Extensions for Segment Routing*

draft-ietf-pce-stateful-pce-14, *PCEP Extensions for Stateful PCE*

RFC 5440, *Path Computation Element (PCE) Communication Protocol (PCEP)*

## Point-to-Point Protocol (PPP)

RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*

RFC 1377, *The PPP OSI Network Layer Control Protocol (OSINLCP)*  
RFC 1661, *The Point-to-Point Protocol (PPP)*  
RFC 1662, *PPP in HDLC-like Framing*  
RFC 1877, *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*  
RFC 1989, *PPP Link Quality Monitoring*  
RFC 1990, *The PPP Multilink Protocol (MP)*  
RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*  
RFC 2153, *PPP Vendor Extensions*  
RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*  
RFC 2615, *PPP over SONET/SDH*  
RFC 2686, *The Multi-Class Extension to Multi-Link PPP*  
RFC 2878, *PPP Bridging Control Protocol (BCP)*  
RFC 4638, *Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE)*  
RFC 5072, *IP Version 6 over PPP*

## **Policy Management and Credit Control**

3GPP TS 29.212 Release 11, *Policy and Charging Control (PCC); Reference points (Gx support as it applies to wireline environment (BNG))*  
RFC 3588, *Diameter Base Protocol*  
RFC 4006, *Diameter Credit-Control Application*

## **Pseudowire**

draft-ietf-l2vpn-vpws-iw-oam-04, *OAM Procedures for VPWS Interworking*  
MFA Forum 12.0.0, *Multiservice Interworking - Ethernet over MPLS*  
MFA Forum 13.0.0, *Fault Management for Multiservice Interworking v1.0*  
MFA Forum 16.0.0, *Multiservice Interworking - IP over MPLS*  
MFA Forum 9.0.0, *The Use of Virtual trunks for ATM/MPLS Control Plane Interworking*  
RFC 3916, *Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)*  
RFC 3985, *Pseudo Wire Emulation Edge-to-Edge (PWE3)*  
RFC 4385, *Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN*  
RFC 4446, *IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)*

---

RFC 4447, *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)*

RFC 4448, *Encapsulation Methods for Transport of Ethernet over MPLS Networks*

RFC 4619, *Encapsulation Methods for Transport of Frame Relay over Multiprotocol Label Switching (MPLS) Networks*

RFC 4717, *Encapsulation Methods for Transport Asynchronous Transfer Mode (ATM) over MPLS Networks*

RFC 4816, *Pseudowire Emulation Edge-to-Edge (PWE3) Asynchronous Transfer Mode (ATM) Transparent Cell Transport Service*

RFC 5085, *Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires*

RFC 5659, *An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge*

RFC 5885, *Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)*

RFC 6073, *Segmented Pseudowire*

RFC 6310, *Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping*

RFC 6391, *Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network*

RFC 6575, *Address Resolution Protocol (ARP) Mediation for IP Interworking of Layer 2 VPNs*

RFC 6718, *Pseudowire Redundancy*

RFC 6829, *Label Switched Path (LSP) Ping for Pseudowire Forwarding Equivalence Classes (FECs) Advertised over IPv6*

RFC 6870, *Pseudowire Preferential Forwarding Status bit*

RFC 7023, *MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking*

RFC 7267, *Dynamic Placement of Multi-Segment Pseudowires*

## **Quality of Service (QoS)**

RFC 2430, *A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)*

RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*

RFC 2598, *An Expedited Forwarding PHB*

RFC 3140, *Per Hop Behavior Identification Codes*

RFC 3260, *New Terminology and Clarifications for Diffserv*

---

## Remote Authentication Dial In User Service (RADIUS)

RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*  
RFC 2866, *RADIUS Accounting*  
RFC 2867, *RADIUS Accounting Modifications for Tunnel Protocol Support*  
RFC 2868, *RADIUS Attributes for Tunnel Protocol Support*  
RFC 2869, *RADIUS Extensions*  
RFC 3162, *RADIUS and IPv6*  
RFC 4818, *RADIUS Delegated-IPv6-Prefix Attribute*  
RFC 5176, *Dynamic Authorization Extensions to RADIUS*  
RFC 6911, *RADIUS attributes for IPv6 Access Networks*  
RFC 6929, *Remote Authentication Dial-In User Service (RADIUS) Protocol Extensions*

## Resource Reservation Protocol — Traffic Engineering (RSVP-TE)

draft-newton-mpls-te-dynamic-overbooking-00, *A Diffserv-TE Implementation Model to dynamically change booking factors during failure events*  
RFC 2702, *Requirements for Traffic Engineering over MPLS*  
RFC 2747, *RSVP Cryptographic Authentication*  
RFC 2961, *RSVP Refresh Overhead Reduction Extensions*  
RFC 3097, *RSVP Cryptographic Authentication -- Updated Message Type Value*  
RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*  
RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions (IF\_ID RSVP\_HOP object with unnumbered interfaces and RSVP-TE graceful restart helper procedures)*  
RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*  
RFC 3564, *Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering*  
RFC 3906, *Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels*  
RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*  
RFC 4124, *Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering*  
RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*

RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*  
RFC 4561, *Definition of a Record Route Object (RRO) Node-Id Sub-Object*  
RFC 4875, *Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)*  
RFC 4950, *ICMP Extensions for Multiprotocol Label Switching*  
RFC 5151, *Inter-Domain MPLS and GMPLS Traffic Engineering -- Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions*  
RFC 5712, *MPLS Traffic Engineering Soft Preemption*  
RFC 5817, *Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks*

## **Routing Information Protocol (RIP)**

RFC 1058, *Routing Information Protocol*  
RFC 2080, *RIPng for IPv6*  
RFC 2082, *RIP-2 MD5 Authentication*  
RFC 2453, *RIP Version 2*

## **Segment Routing (SR)**

draft-francois-rtgwg-segment-routing-ti-lfa-04, *Topology Independent Fast Reroute using Segment Routing*  
draft-gredler-idr-bgp-ls-segment-routing-ext-03, *BGP Link-State extensions for Segment Routing*  
draft-ietf-isis-segment-routing-extensions-04, *IS-IS Extensions for Segment Routing*  
draft-ietf-mpls-spring-lsp-ping-02, *Label Switched Path (LSP) Ping/Trace for Segment Routing Networks Using MPLS Dataplane*  
draft-ietf-ospf-segment-routing-extensions-04, *OSPF Extensions for Segment Routing*

## **Synchronous Optical Networking (SONET)/Synchronous Digital Hierarchy (SDH)**

ANSI T1.105.03, *Jitter Network Interfaces*  
ANSI T1.105.06, *Physical Layer Specifications*  
ANSI T1.105.09, *Network Timing and Synchronization*  
ITU-T G.703, *Physical/electrical characteristics of hierarchical digital interfaces*  
ITU-T G.707, *Network node interface for the synchronous digital hierarchy (SDH)*  
ITU-T G.813, *Timing characteristics of SDH equipment slave clocks (SEC)*

- ITU-T G.823, *The control of jitter and wander within digital networks which are based on the 2048 kbit/s hierarchy*
- ITU-T G.824, *The control of jitter and wander within digital networks which are based on the 1544 kbit/s hierarchy*
- ITU-T G.825, *The control of jitter and wander within digital networks which are based on the synchronous digital hierarchy (SDH)*
- ITU-T G.841, *Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum 1, issued in July 2002*
- ITU-T G.957, *Optical interfaces for equipments and systems relating to the synchronous digital hierarchy*

## **Time Division Multiplexing (TDM)**

- ANSI T1.403, *DS1 Metallic Interface Specification*
- ANSI T1.404, *DS3 Metallic Interface Specification*

## **Timing**

- GR-1244-CORE, *Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005*
- GR-253-CORE, *SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000*
- IEEE 1588-2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*
- ITU-T G.781, *Synchronization layer functions, issued 09/2008*
- ITU-T G.813, *Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003*
- ITU-T G.8261, *Timing and synchronization aspects in packet networks, issued 04/2008*
- ITU-T G.8262, *Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007*
- ITU-T G.8264, *Distribution of timing information through packet networks, issued 10/2008*
- ITU-T G.8265.1, *Precision time protocol telecom profile for frequency synchronization, issued 10/2010*
- ITU-T G.8275.1, *Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, issued 07/2014*
- RFC 5905, *Network Time Protocol Version 4: Protocol and Algorithms Specification*

---

## Two-Way Active Measurement Protocol (TWAMP)

RFC 5357, *A Two-Way Active Measurement Protocol (TWAMP) (server, unauthenticated mode)*

RFC 5938, *Individual Session Control Feature for the Two-Way Active Measurement Protocol (TWAMP)*

RFC 6038, *Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features*

## Virtual Private LAN Service (VPLS)

RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*

RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*

RFC 5501, *Requirements for Multicast Support in Virtual Private LAN Services*

RFC 6074, *Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)*

RFC 7041, *Extensions to the Virtual Private LAN Service (VPLS) Provider Edge (PE) Model for Provider Backbone Bridging*

RFC 7117, *Multicast in Virtual Private LAN Service (VPLS)*

## Voice and Video

DVB BlueBook A86, *Transport of MPEG-2 TS Based DVB Services over IP Based Networks*

ETSI TS 101 329-5 Annex E, *QoS Measurement for VoIP - Method for determining an Equipment Impairment Factor using Passive Monitoring*

ITU-T G.1020 Appendix I, *Performance Parameter Definitions for Quality of Speech and other Voiceband Applications Utilizing IP Networks - Mean Absolute Packet Delay Variation & Markov Models*

ITU-T G.107, *The E Model - A computational model for use in planning*

ITU-T P.564, *Conformance testing for voice over IP transmission quality assessment models*

RFC 3550 Appendix A.8, *RTP: A Transport Protocol for Real-Time Applications (estimating the interarrival jitter)*

RFC 4585, *Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)*

RFC 4588, *RTP Retransmission Payload Format*



## **Wireless Local Area Network (WLAN) Gateway**

3GPP TS 23.402, *Architecture enhancements for non-3GPP accesses* (S2a roaming based on GPRS)



# Customer Document and Product Support



## Customer Documentation

[Customer Documentation Welcome Page](#)



## Technical Support

[Product Support Portal](#)



## Documentation Feedback

[Customer Documentation Feedback](#)

